Chapter 11

Why Pension Fund Management Needs a Paradigm Shift

Keith Ambachtsheer

A paradigm is a lens through which to see the world. The lens through which the institutional investment community has heretofore seen the defined benefit (DB) pension fund management world has three facets. First, stocks have outperformed bonds by about 5 percent per annum on average in the past. Looking ahead, many DB plan managers believe that this 5 percent equity risk premium will also be available in the future, as long as one is patient enough. Second, many DB plan managers believe that a 60:40 equity-bond mix provides sufficient diversification over the shorter term to create long-term sustainability, while still providing a 3 percent risk premium as compared to a 100 percent bonds-only policy. They see this as a good reward/risk 'deal'; in fact, the experience of the 1990s persuaded many that even a more aggressive 70:30 mix could pass the prudence test. Given the first two assumptions, resources allocated to managing pension funds can mainly focus on generating additional net return relative to a passively implemented 60:40 (or 70:30) policy portfolio.

The fundamental questions we address in this chapter flow from these assumptions. Is this traditional lens a good one through which to see the pension fund management world in the future, and if not, why not? What would a better lens impart? Our approach is inductive, beginning with a thorough examination of 'old' paradigm pension fund management practices and the results they have produced. We unearth material flaws in current practices and suggest better ways to manage pension funds.

Performance under the Old Paradigm

It is well known that what gets measured, gets managed; hence it is not surprising that the measurement of pension fund performance has historically focused on the third facet of the 'old' paradigm enunciated above. Specifically, managers have been alert to how the fund has performed relative to simply implementing the chosen asset mix policy passively. In other words, managers ask was additional return produced, and at what cost and risk? Answers to this question are available from the Cost Effectiveness Measurement (CEM) database containing 10 years of pension

performance data.[1] The database includes information on 256 funds (142 American [$1.5 trillion], 90 Canadian [C$374 billion], and 12 each European [€328 billion] and Australian [A$70 billion]); of these, 114 were corporate, 113 public sector, and 29 industry/other funds. Average fund size was $9 billion versus a median size of $2 billion.

In the reported results that follow, all available performance data over the ten-year period was employed including information on funds with less than a full ten-year history. Specifically, all possible calculations were made on the measure equal to [(annual total fund return) − (annual fund policy portfolio return) − (annual total operating costs)]. The resulting 2,671 fund 'net implementation value-added' (NIVA) metrics were analyzed. Figure 11-1 reveals the results: the average NIVA was +11 basis points (bps), with a respectable *t*-value of 1.8, suggesting that active management did add value over the 1993–2002 period. Even more interesting is the split of total NIVA into its three components. Average 'gross in-category value-added' was an even more impressive +69 bps ($t = 11.2$), while 'mix value-added' subtracted an average −21 bps ($t = -5.9$). Operating costs brought total fund performance down a further average −37 bps.

These findings lead to two further important conclusions. First, operating with a 100 percent passive implementation style costs money, so the proper benchmark against which to evaluate the +11 bps NIVA finding is not 0 bps, but the cost of pure passive implementation including all relevant overheads. This consideration reduces the benchmark NIVA from 0 basis points to, say, minus 9 bps. Consequently, a more realistic estimate of average NIVA experience of the pension funds in the CEM database, relative to a pure passive alternative in the 1993–2002 period, was +20 basis points. Second, the observed average NIVA of +20 bps would have more than doubled, had it not been for the calculated negative ten-year 'mix' effect of −21 bps. This effect resulted from the funds doing something other than following the standard database protocol of rebalancing the asset mix weights back to those of the stated asset mix policy on January first of each year. Below, we address the question of why actual rebalancing strategies systematically underperformed, as compared to the standard rebalancing protocol over the period.

The Asset Class 'Alphas': Good News and Bad News

Also worth investigating further is the question of which asset classes produced the greatest 'in category' net value added, and whether we can identify structural pension fund characteristics that are statistically associated with positive NIVA performance. Figure 11-2 displays the 'in-category selection alphas' for seven major asset classes for all US funds in the CEM database. These 'in-category alphas' are calculated by subtracting from the annual asset class return, the return on a relevant asset class benchmark

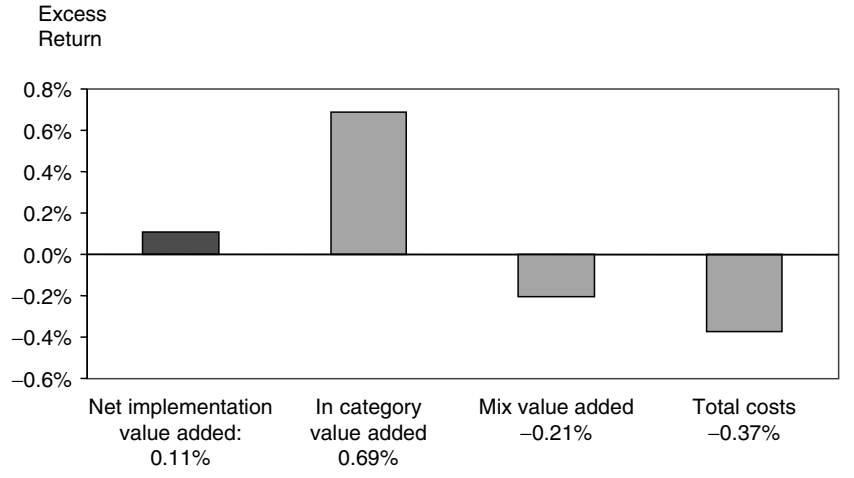

Figure 11-1. Did active management add value?
Source: Cost-Effectiveness Measurement Inc.

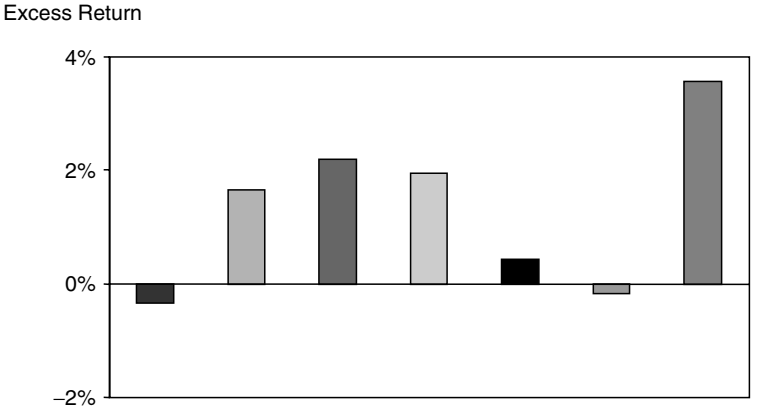

	Stock: Domestic large	Stock: Domestic small	Stock: Foreign	Stock: Emerging	Fixed income	Real estate	Private equity
Average net value added (%)	–0.3	1.7	2.2	1.9	0.4	–0.2	3.6
Standard deviation (%)	6.5	13.2	7.6	9.1	3.6	7.8	33.7

Figure 11-2. Actively managed asset class alphas for US funds.
Source: Cost-Effectiveness Measurement Inc.

portfolio selected by the participating fund. Note also that the direct costs associated with managing investments within each of the seven asset classes (i.e. either external fees or allocated internal costs) have been netted out.

In evaluating Figure 11-2, one should be mindful that the averages and standard deviations were calculated from very large actively managed samples over the 1993–2002 period. For example, the domestic large cap stocks category statistics are drawn from a sample of 1,770 annual observations. Even the real estate and private equity sample sizes are 1,466 and 1,221 respectively. So from the perspective of standard tests of statistical significance, all the calculated average asset class net alphas easily pass the null hypothesis test, except for real estate. That leaves us with five positive asset class alpha results, and only one negative one. That is the good news.

Having said that, the results also confirm that individual fund asset class alpha outcomes vary greatly: this is the bad news. From a reward/risk ratio perspective, none of the average asset class alpha performances are exciting. The attractiveness of average selectivity within foreign stock investing in developed markets scores best with an average reward/risk ratio of 0.3 (i.e. 2.2/7.6), and we suspect this success largely represents the systemic under-weighting of Japan by most US pension funds over the 1993–2002 period. The very modest average 0.1 reward/risk ratio generated by the private equity asset class provides a very sober view of selectivity risk here. The 33.7 percent standard deviation suggests that being right on the asset class is not enough. It means that, even though private equity as an asset class may perform well, one can still end up with terrible results in the fund. The implication is that, unless a fund has enormous confidence in its own private equity manager selection skills, it had better pay careful attention to manager diversification.

Of course, the observed great variance in private equity results could be given a more positive spin. The results suggest that good manager selection skills in this asset class have a very large payoff. Figure 11-2 statistics suggest that similar conclusions (both negative and positive) apply to the domestic small cap stocks and foreign emerging markets stocks asset classes.

The Negative Mix-Related Alpha Question: Is 'Excess Cash' the Culprit?

We next focus on the highly statistically significant −21 bps average annual loss due to asset mix-related reasons over the ten-year period displayed earlier in Figure 11-1. Further research points to a likely contributor: in the eight individual years when average fund policy returns were positive (i.e. 1993–2000), the average asset mix–related alphas were negative in each year. In contrast, in the two individual years when the average fund policy returns were negative (i.e. 2001 and 2002), the average asset mix–related alphas were positive in both years.

What does this year-by-year breakdown of the ten-year average −21 bps result imply? It suggests that the holding of excess cash in relation to the policy asset mix may be a culprit, a hypothesis consistent with the average asset mix–related alpha being negative in each of the eight positive policy return-years (average mix alpha was −32 bps), and being positive in each of the two negative policy return years (average mix alpha was +28 bps). The discovery that holding 'excess cash' may reduce the typical fund's incremental return by one-third sends a powerful message. Decisive steps should be taken to minimize the negative impact of excess cash on total fund long-term performance.

Is There Also a Pay-off from Actively Managing Pension Fund Costs?

Figure 11-1 indicated that investment-related total operating costs captured in the database averaged 37 bps annually over the ten-year period. Table 11-1 provides important additional information, where we note that the database spans annual cost experience from a miniscule 1 bp to a hefty 236 bps (based on 2,450 observations). Indeed, the 10–90 percentile experience still encompasses a broad 15–58 bps range. This wide-ranging operating cost experience prompts an important question: is there an 'excess cost' performance drag, in the same way that the analysis above suggested there was an excess cash performance drag, over most of the 1993–2002 period?

The simplest way to address this question is to regress fund performance (NIVA) against total operating costs. If there was an excess cost performance drag embedded in the data, it would show up as a statistically significant negative cost coefficient—a hypothesis that seems to be the case. The statistically significant cost coefficient was −0.48 ($t = -2.2$), implying that one basis point of additional operating cost reduced NIVA by an average −0.48 bps over the 1993–2002 period. While this is an important finding, it raises three important further questions. First, pension fund management

Table 11-1 Annual Total Operating Costs (bps)

Low	1.1
10th percentile	14.5
1st quartile	23.1
Median	33.9
3rd quartile	46.7
90th percentile	58.4
High	236.2

Source: Cost-Effectiveness Measurement Inc.

should experience significant economies of scale; what happens when we adjust for this scale factor? Second, certain types of investment policies are much more expensive to implement than others; what happens when we adjust for this investment policy factor? Third, operating cost structures may to some degree be country-specific; what happens when we adjust for this country factor?

All three questions have interesting answers. First we find that the scale effect averaged −19 bps. For every tenfold increase in fund asset value in the database (total range $100 million–$150 billion), total operating costs declined an average −19 bps. Thus large funds have a material cost advantage over small funds. Second, the incremental cost coefficients for domestic stock, foreign stock, real estate, and private equity mandates were +38, +62, +88, and +199 bps. These metrics indicate average total operating cost increases by shifting a fund out of bonds into domestic stocks, foreign stocks, real estate, and or private equity respectively. Thus investment policy choices do materially impact cost experience. Finally, adjusted for scale and mix factors, Canadian funds experienced operating costs an average −11 bps lower than their US counterparts. It appears that Canadian funds generally chose simpler investment management structures and faced lower domestic external management fee structures. Out of the total cost variance (measured by a standard deviation of 19 bps for the full sample), these three cost drivers explained about one-half of the cost variance.

Next we examine whether the three systematic cost drivers are responsible for the measured 'excess cost' performance drag coefficient of −0.48 bps. This is assessed by reestimating the fund NIVA performance versus fund operating costs relationship, but now using fund cost experience adjusted for differences in scale, asset mix, and country. If the 'excess cost' coefficient remains statistically significantly negative, it would indicate that—regardless of fund size, mix, or country—eliminating excess costs, like eliminating excess cash, has a positive performance payoff. The new excess cost coefficient is, in fact, negative and significant [−0.54; $t = -1.9$]. Thus every basis point of additional scale-mix-country-adjusted excess cost reduced NIVA by an average −0.54 bps over the 1993–2002 period. This suggests that managing costs by ensuring that they have a positive payoff matters, which is the essence of cost-effectiveness.

Two Cost-Effectiveness Strategies

What might specific examples of such 'cost-effectiveness' strategies be? The CEM database offers two examples. First, there is no statistical difference between the gross alphas of comparable externally and internally managed investment mandates. However, internal management is on average 30 bps less costly and hence that strategy outperformed external

management by an average 30 bps on a net basis. Second, every 10 percentage point increase in passive management was statistically associated with nine bps of additional NIVA. This suggests that significant exposure to very low-cost passive management, together with concentrated active management, was generally a cost-effective, value-producing combination over the period.

These findings add up to two important lessons. In assessing the attractiveness of within-asset class active management, it is first important to distinguish between average experience, and the variability of that experience. For example, over the 1993–2002 period, active management was on average successful in five out of seven major asset classes. However, with that average success came significant potential for individual funds to reduce return rather than enhance it. Second, a much surer way to increase fund NIVA is to carefully manage excess cash and excess costs. The value of having neither was easily worth as much as the total average value of active management, but without its downside exposure, suggesting that there may, after all, be something akin to the proverbial 'free lunch'.

The Common Performance Driver: Organizational Effectiveness

Our findings in turn beget broader questions. For instance, can we identify a common driver of good (or bad) 'alpha' performance? What can we learn about the attributes of high performance pension funds? Research during the 1990s provides interesting answers to these questions. For instance, anthropologists William O'Barr and John Conley (1992) caused quite a stir with their book on pension fund management drawn from close observations of nine pension fund organizations over the 1990–1 period. These evaluations led them to conclude that fund organizational structures were the result of the funds' historical origins, rather than from the disciplined application of good governance and management practices. They also found that deflecting responsibility and shifting blame seemed to be important motivators in the funds' decision-making processes. For instance, close relationships with outside services suppliers (e.g. money managers, consultants, investment dealers, custodians) seemed to be an extraordinarily high priority for pension fund trustees and executives.

The pension community of the day reacted with shock and outrage to these allegations, but when, two years later, fifty pension fund CEOs were asked what they estimated their 'excellence shortfall' to be (i.e. annual return shortfall due to internal governance and management problems), the median response was a hefty 66 bps. When they characterized the nature of their internal problems, most mentioned poor structure and decision processes (forty-nine mentions), almost half mentioned inadequate resources (twenty-four mentions), and almost half spoke of lack of

focus and mission clarity (twenty-two mentions). This feedback suggested that perhaps O'Barr and Conley had found factors that did actually make a difference.

Logically then, the next step in the performance driver discovery process was to establish a statistical relationship between pension fund performance and the quality of governance and management practices. In an initial study, Ambachtsheer et al. (1998) used 1993–6 NIVAs from the CEM regressed on metrics representing the quality of a pension fund's governance and management practices. These were calculated from the survey responses of eighty participating pension fund CEOs. The heart of the survey was forty-five statements related to a fund's governance, management, and operations practices. The CEOs were invited to give each of the statements a ranking between one and six, depending on their assessment how well the statement reflected reality inside their own organization. Each statement was designed so that a six represented perceived best practice and a one perceived poor practice. In each case, the average of the forty-five statement rankings in each survey was designated as the 'CEO Score' for that survey. Thus in all, eighty 'CEO Scores' were calculated from the eighty completed surveys.

The eighty CEO Scores were interesting: for instance, scores ranged from a low of 3.0 to a high of 6.0, with both mean and median scores at 4.8. This meant that overall, the eighty raters felt pretty good about their organizations (not one was below three at the low end, and one wag assigned his organization a perfect six on forty-five counts). Having said that, an overall three to six score spread is a good statistical range to work with. Also, of the eighty CEOs, fifty-one represented corporate funds and twenty-nine noncorporate (mainly public sector). The average corporate CEO Score was 4.9 versus a 4.7 average for noncorporates; the difference was not statistically significant though it suggested that corporate fund CEOs had, on average, a somewhat higher opinion of their fund's organizational competence than did noncorporate fund CEOs of theirs. This finding fits nicely with our earlier observation that corporate funds had somewhat higher average NIVA performance over the 1993–2002 period (12 bps) than did noncorporate funds. The statistical relationship between the CEO Scores and fund investment performance, the data indicated a 'CEO Score' coefficient of 0.4 with a significant t-value of 2.9, using 1993–6 performance data. This meant that an additional point of CEO Score in 1997 translated into an additional average 40 bps of annual return over the 1993–6 period. A two-point 'CEO Score' increase, say from 3.5 to 5.5, was associated with 80 bps of annual incremental performance. This finding fitted nicely with the median 66 bps excellence shortfall estimate of the fifty CEOs at the 1994 New York gathering.

We have now updated the analysis to 2000, as reported in Table 11-2. Arguably, finding a statistical relationship between the CEO Scores and

TABLE 11-2 New Results: Fund Performance vs. 'CEO Scores'

	'CEO score' coefficient	*t-value*
Pre-survey period		
4-year period (1993–6)	+0.4	2.9
Post-survey period		
1-year period (1997)	+0.8	2.7
2-year period (1997–8)	+0.6	1.9
3-year period (1997–9)	+0.7	1.7
4-year period (1997–2000)	+0.3	1.0

fund performance in this post-survey period would provide further confirmation that good governance and management matter most. First we replicate previous findings for the 1993–6 period, and then we set out the new results for 1997 only, then for 1997–8, then for 1997–9, and finally for the full 1997–2000 period.

The evidence shows that funds with high CEO Scores were able to maintain their average performance edge, relative to funds with low CEO Scores, even after immediately after the survey was conducted in 1997. It was only in the fourth year after the rankings were completed that there was a material reduction in the performance edge of high CEO Score funds relative to the low scoring funds. Note that the post-survey period coefficients were higher than the pre-survey coefficient of 0.4, until the year 2000 was included in the performance results. So it appears that a careful study of governance and management practices can help in understanding and hence improving pension fund performance.

Good Governance and Management Related to NIVA Performance

We should recall that the 'CEO Scores' were composites of forty-five individual statement rankings: approximately one-third of these statements related to governance practices, one-third to management practices, and one-third to operations. A test to determine which of the forty-five statement scores correlated most closely with fund performance showed that only eleven had significant positive correlation with fund performance. Of these, six were governance-related and five management-related.

The statistically significant governance statements in the survey related mainly to effective fiduciary behavior and selection processes, clarity in delegation of authority, and a high level of trust between the governing and managing fiduciaries. The statistically significant management statements related mainly to clear strategic positioning and to the effective

development and execution of the fund's strategic plan. These findings represent a noteworthy convergence between what are deemed to be generally good governance and management practices in the for-profit and not-for-profit sectors as a whole, and what we now find is important in the governance and management of pension funds.

An interesting issue is why the performance edge of the high CEO Score funds began to diminish in the fourth year after the scores were captured. Perhaps the high scoring funds of 1997 failed to work hard enough to maintain a competitive advantage beyond 1997. Certainly, we find that maintaining a competitive advantage requires as much work as it took to achieve that advantage in the first place. This question deserves further research.

Assessing Pension Fund Risk

Until this point, the discussion has not considered the risk side of the pension fund management task. Figure 11-3 plots the net alphas (NIVAs)

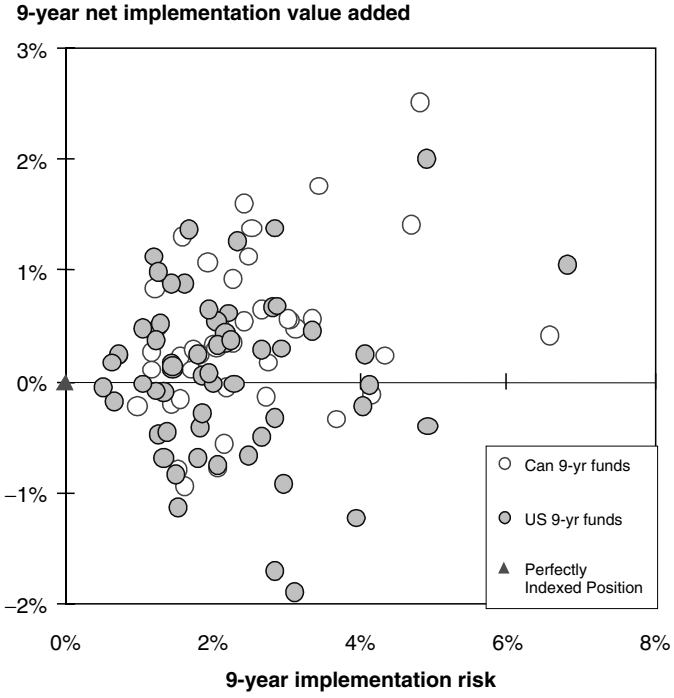

Figure 11-3. Implementation value added vs. implementation risk.
Source: Cost-Effectiveness Measurement Inc.

of all eighty US and Canadian funds with ten years of continuous data against their implementation risk (i.e. the volatility of fund returns relative to their respective policy portfolio returns). This implementation risk can be thought of as the additional risk contributed by active management. As noted above, the average NIVA realization was marginally positive. Now we also see that the measured implementation risk for these funds ranged from a low of virtually zero, to a couple of outliers in the six percent range. Nevertheless, the bulk of measured experience was around two percent. Statistically, the NIVA/Tracking Error Volatility coefficient was a positive 0.12, with a *t*-value of 1.9. So on average, funds garnered a marginal amount of additional net alpha for a modest amount of additional 'active' risk. Thus our previous conclusion continues to hold: in an old paradigm context, pension fund management was on average modestly successful over the ten-year period ending 2002.

Next we explore the behavior of the fund return components that can be attributed to the funds' choices of asset mix policy. In the CEM data file, these choices are represented by passively implemented 'policy portfolios'. Specifically, we are interested in the behavior of these policy portfolios relative to the liability return requirements of each fund.[2] These requirements are estimated in each case by calculating the return on a bond portfolio (i.e. the 'liability portfolio') that mimics a fund's liabilities in terms of duration and inflation-sensitivity. Figure 11-4 addresses this ques-

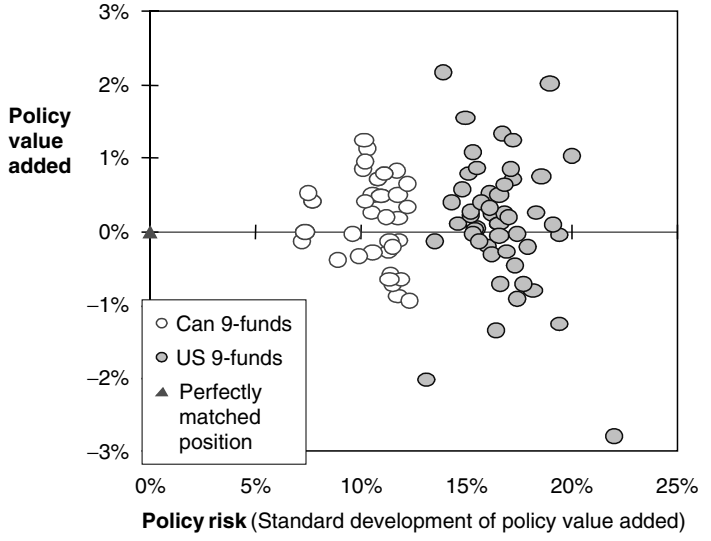

Figure 11-4. Policy value added vs policy risk.
Source: Cost-Effectiveness Measurement Inc.

tion and it tells a much less positive story. In comparison with Figure 11-3, for instance, nothing much changes on the vertical axis. Over the observation period, policy portfolio returns over/underperformed their respective 'liability portfolio' returns in a narrow +2% to −2% range (i.e. Policy Value Added). However, looking along the horizontal axis (i.e. Policy Risk), one notes that the policy mismatch volatility metrics for the Canadian funds are now in the 7–12 percent range. For the US funds, the range is more like 15–20 percent. Further, there is no correlation between the liability-relative excess returns and liability relative policy risks. So we derive two important conclusions: first, the 'policy' risk of US funds dominated their 'active' risk by a factor of about eight to one over the ten-year observation period. Second, this outsized exposure to policy risk received, on average, no compensation.

The database also provides direct evidence that, consistent with 'old' paradigm thinking, fund managers of late have made no measurable effort to take into account the financial characteristics of their liabilities when structuring their fund policy portfolios. A test was devised to compare the asset mix policies of the 50 percent of US funds with the most inflation-sensitive liabilities (82 percent inflation-sensitive, on average) with the 50 percent with the least inflation-sensitive liabilities (31 percent inflation-sensitive, on average). It turned out that the two subuniverses had, on average, identical asset mix policies: 63 percent equities, 32 percent fixed income, and 5 percent inflation-sensitive investments. Even if the two subuniverses had the same appetite for equity risk, one would expect that the most inflation-sensitive half of the universe would favour inflation-sensitive investment over fixed income investments. The reverse should be true with the less inflation-sensitive subuniverse. Remarkably, this was not the case.

The bottom line is that, when both recent bull market and bear market experiences are included in the measurement period, and when we benchmark the realized policy returns against marked-to-market liability portfolio requirements, the policy portfolios of DB pension funds generated very little excess return for a highly material amount of balance sheet mismatch risk in the ten years ending in 2002. (Thus, Figure 11-4 showed that most US funds had policy risks in the 15–20 percent range.) Of course under 'old' paradigm behavior, these findings should not be surprising. We know, for example, that 'old' paradigm typically behavior produces very similar policy asset mixes without regard for differences in liability structures. We also know that the observation period included the three-year 'perfect pension storm' period (i.e. 2000–2), when both equity prices and bond yields fell together; this led to massive marked-to-market funded pension ratio declines in the 25–50 percentage point range. The profound question raised by these findings is whether the 'old' fund management paradigm indeed is the best lens through which pension fiduciaries should see their world.

Why the 'Old' Lens Distorts

The old lens is not the best lens for pension fiduciaries to use today for at least three reasons. First, the prospective equity risk premium will not always be 5 percent, even over long periods: it can be predictably high, normal, or low, with 'normal' about 2.5 percent rather than 5 percent. Second, a 60:40 (or 70:30) asset mix policy is not always a good reward/risk 'deal' for the stakeholders in pension plans. This is because such a policy will not always offer sufficient reward per unit of properly defined risk; further, in some situations, such a policy is just absolutely too risky, regardless of whether the prospective reward is sufficient or not. Third, devoting the bulk of a fund's resources to attempt to generate a modest amount of policy portfolio-relative 'net alpha' for a modest amount of policy portfolio-relative additional risk may no longer be prudent.

To confirm that the equity risk premium is indeed predictably high, medium, or low, we point to Figure 11-5 which plots almost 200 years of equity risk premium predictions on the horizontal axis, against on the vertical axis what actually happened (i.e. the actual excess return of equities

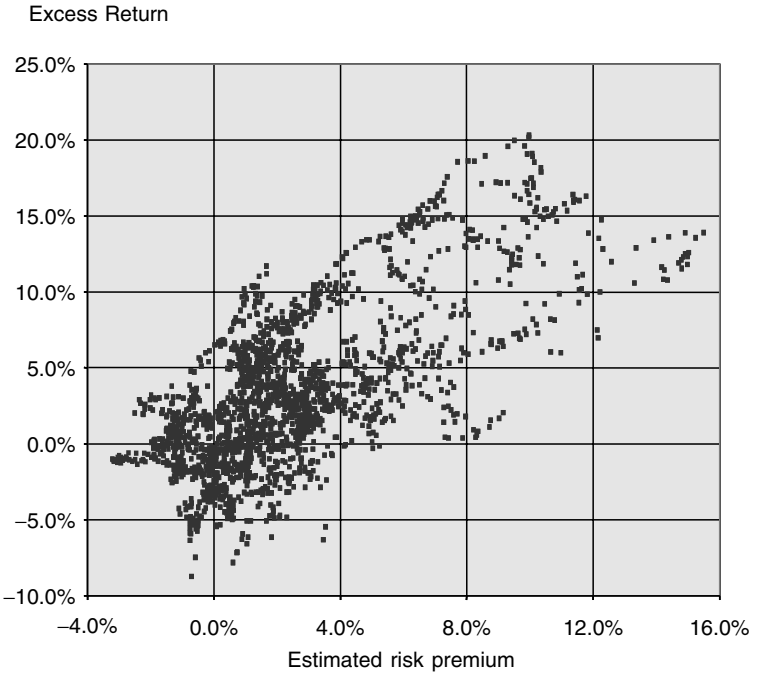

Figure 11-5. Equity risk premium & subsequent ten-year excess returns, 1810–1991. *Source*: Arnott and Bernstein (2002).

relative to bonds) over the ten-year period after the predictions were made. The figure uses the simple rule of thumb that the ERP at any point in time equals the current dividend yield plus the prior ten years earnings growth minus the current long bond yield. Note the predictions on the horizontal axis range from minus three percent to +15 percent. The actual ten-year outcomes on the vertical axis range from minus eight percent to +20 percent (Arnott and Bernstein 2002).

The key issue is whether there is a positive relationship between the equity risk premium predictions and the subsequent ten-year outcomes purely random. The data show that this was indeed the case. Generally, high predictions produced high outcomes, and low predictions low outcomes. The rule of thumb equity risk premium prediction calculation today works out to about +1.5 percent.[3] The figure says that when the prediction is +1.5 percent, the historical ten-year outcome range has been about −7 percent to +10 percent. There are no good reasons that we know of why this range should not also frame ten-year equity-bond return prospects today. If it does, the case for placing most of the risk chips in pension funds on a policy portfolio with 60–70 percent equity exposure is questionable.

Figure 11-5 also makes clear that even if the rule of thumb produced an expected equity risk premium today higher than 1.5 percent, a policy portfolio with 60–70 percent equity exposure might still be too risky, in some situations. Pension funds of DB plans sponsored by corporations struggling with poor operating results and weak balance sheets come to mind. The figure also shows that in a world of varying equity risk premium prospects, and varying abilities of fund stakeholders to withstand disappointing outcomes in five to ten year timeframes, focusing largely on alpha games, where the realistic stakes are a net 0.5 percent of extra return for two percent of extra risk, makes little sense.

A New Lens

So pension fund managers need a new lens. Here are its three facets. First, *a set of investment beliefs* grounded in good theory and confirmed by real world experience. An important element of such a beliefs set should be the notion of a predictably varying equity risk premium. Other elements will reflect views on the return generation processes in various capital market sectors and the predictability of those processes. Second, *an integrative investment model* is needed that directly links a varying return opportunity set to stakeholder income needs and risk tolerances. The essence of a successful model is to capture the linkages between fuzzy expectations, transactional friction, and stakeholder needs. Third, *a decision-making protocol* must be designed that can dynamically integrate components one and two into 'value' for fund stakeholders. The essence of a successful protocol

is effective human interaction. Thus its foundation must be alignment of economic interests, good governance, and good organization design.

The implications of seeing pension fund management through this new three-component lens can be best understood by being mindful of Einstein's observation that solving significant problems requires a higher level of thinking than used to create them. For example, while Keynes' 1936 opus *The General Theory* was meant to address the economics of depression, it contains a very powerful chapter on investment beliefs (Keynes 1961 Chapter 12). With some restatement, Keynes believed that capitalism embodies two distinct types of investment processes. 'Long horizon' investing is about the projection and valuation of future cash flows. 'Short horizon' investing is about zero-sum, adversarial trading in financial instruments. The mistake many pension funds make is to try and blend these two radically different investment processes together. In the 'old' paradigm, 'long horizon' strategies are assumed to have pre-determined long term, static expected rates of return. 'Short horizon' strategies, on the other hand, require 'smart' active managers capable of producing positive net alpha. So, for example, most 'active' investment mandates today continue to be defined in the context of static 'long horizon' benchmark portfolios, to which active managers are to add some 'alpha' through over/underweighting strategies.

This is a highly restrictive, suboptimal delegation strategy, if two plausible conditions hold. The first condition is that the expected returns on 'long horizon' investment strategies are not static but variable, and to some degree predictive. The second condition is that successful 'active' strategies can be identified before the fact. These two conditions logically lead to the internally controlled, dynamic, integrated 'new' paradigm reward/risk management process further described below. Why is the long horizon/short horizon distinction so important in this? Because there is persuasive evidence that the returns on 'long horizon' investments are intrinsically predictive, to some degree. In contrast, whether the returns on 'short horizon' investment processes (most efficiently packaged as bundles of active market-neutral strategies) are predictive, is in the eyes of the beholder. But what may be true for some can't possibly be true for all.

There is no place for static 'policy portfolios' in this 'new' investment paradigm. The integrative investment model sketched out in Figure 11-6 makes clear why. The only relevant benchmark now is the risk-minimizing portfolio that looks most like the liabilities the pension fund is meant to cover. Moving to the right on the horizontal axis implies taking on increasing mismatch risk relative to this liability portfolio. In the 'new' paradigm, all mismatch risk is 'active' risk. It no longer matters whether its source is 'long horizon' risky strategies, 'short horizon' risky strategies, or some blend of the two. The vertical axis measures increasing amounts of expected or realized excess return relative to the return of the risk-minimizing portfolio. The 10 percent on the horizontal axis might

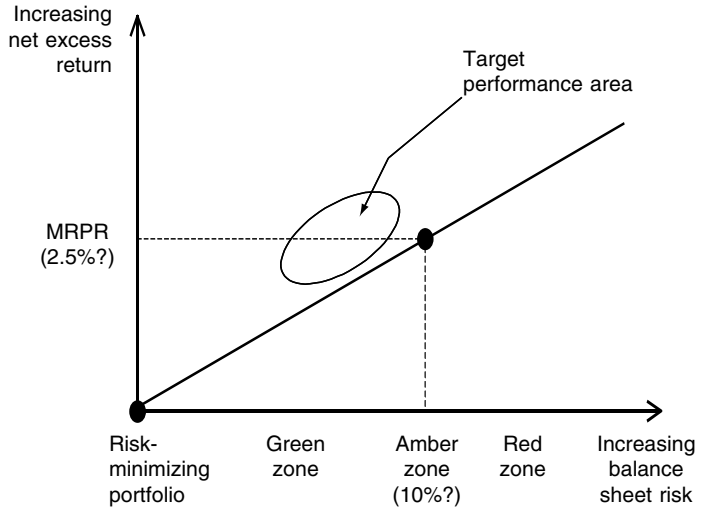

Figure 11-6. Integrative investment model.

represent the 'mismatch risk budget' that a fund management must work within. The 1.5 percent on the vertical axis might represent the minimum required price of risk (MRPR) that should be earned with a 10 percent risk budget. The ellipse indicating the target performance area suggests that now all that matters are balance sheet–relative results, and not policy portfolio–relative results. Thus that is also what should be measured.

Earlier in the chapter, we provided statistical evidence that good governance and organization design matter. Such human system-based disciplines lead to mission clarity, delegation clarity, effective strategic planning and execution, and a high level trust within the pension fund organization. These factors in turn drive superior organization performance, and getting them right is just as important as getting the investment beliefs and the investment model right. Specifically then, these factors become important in actually moving the organization from the 'old' to the 'new' paradigm, and in setting up the necessary infrastructure to run a 'new' paradigm pension fund organization.

If this is the new standard against which to evaluate the practices of today's institutional investment community, how does it measure up? Unfortunately, most of the discussions we hear at various institutional investment forums continue to fit the 'old' investment paradigm far better than the new. Nevertheless, there are glimmers of light, since some pension fund executives are observing that they are being pressed by their boards 'to become more tactical'. We take this to mean that some boards are beginning to understand there is no such thing as a single policy portfolio

for all seasons. Instead, the emerging view is that balance sheets should be managed dynamically, sensitive to changing socioeconomic and capital market conditions.

Another encouraging sign is the fact that conference agendas have been giving greater attention to the importance of good fund governance and organization design. Pension boards and investment executives themselves have begun to pay more attention to these organizational factors. Pension consultants are becoming more open to the possibility of a paradigm shift. Investment managers with confidence in their investment skills are beginning to see their world through a lens that splits investment strategies as being either 'long horizon' or 'short horizon' based.

Conclusions

Pension fund management has been guided by a paradigm that appeared to work well during the 1980s and the 1990s. It seemed to be a good lens through which to see the world. The equity risk premium was generally positive, as assumptions suggested it should be. Equity market dips were short, and soon reversed themselves. Nothing happened that a 60–40 asset mix policy apparently couldn't deal with. In addition, our research confirmed that over the ten-year 1993–2002 period, a typical fund generated a modest amount of 'net alpha' by taking a modest amount of 'implementation' risk relative to its policy portfolio. Under the 'old' paradigm, this was an important accomplishment.

Then, during the three-year 'perfect pension storm' of 2000–2, the 'old' lens developed serious cracks. The surpluses on DB balance sheets turned to serious deficits. With the benefit of hindsight, it became clear that the asset mix policies of the 1990s exposed the stakeholders of DB balance sheets to material mismatch risk. This experience has led many pension boards and executives to ask if that material mismatch risk has now receded. Alternatively, could pension balance sheet stakeholders withstand the financial pain if a 2000–2 type shock were to occur again a few years hence. Related issues are how much compensation in terms of additional expected return is sufficient for these risks to be undertaken, and what the sources of that addition return will be in the future. Any pension organization that thinks these questions are important can no longer be guided by the 'old' paradigm. It needs a new lens through which to see the world.

Endnotes

1. Cost-Effectiveness Measurement Inc. (CEM) benchmarks the organizational performance of many of the world's largest DB and DC pension plans, both on the investment and benefit administration sides of the pension 'business'. The author is a Founding Partner of the firm, located in Toronto, Canada.

2. CEM estimates the 'liability returns' of participating DB pension plans by first creating a bond portfolio that reflects the inflation sensitivity and the duration of plan liabilities, and then estimating the marked-to-market annual return on that portfolio.
3. The current stock dividend yield is about 1.5 percent. The referenced Arnott-Bernstein paper (2002) suggests a good dividend real growth rule of thumb is 2 percent. The current yield on long TIPS is about 2 percent. So using the simplified Gordon formula, we are looking at a current equity risk premium estimate of $\{1.5\% + 2\% - 2\% = 1.5\%\}$.

References

Ambachtsheer, Keith, Ronald Capelle and Tom Scheibelhut (1998). 'Improving Pension Fund Performance', *Financial Analysts' Journal*, 54/6:15–20.

Arnott, Robert and Peter L. Bernstein (2002). 'What Risk Premium is Normal?', *Financial Analysts' Journal*, 58/2: 64–85.

Keynes, John Maynard (1961). *The General Theory of Employment, Interest and Money*. New York: St. Martin's Press.

O'Barr, William and John Conley (1992). *Fortune and Folly: The Wealth and Power of Institutional Investing*. Homewood, IL: Business One Irwin.