Contents

Chapter 2

Actuarial Cost Methods

2.1 Why and Wherefore

Cost methods apply to pension plans, so we cannot begin to understand what a cost method is, how it works, or why we need to have one, without saying a little about pension plans in general. Although the term "pension plan" is applied to a bewildering variety of retirement and deferred-compensation schemes, we shall use the term *pension plan* in a very restrictive sense: any arrangement for providing monthly payments to a person for life beginning at a stated age, where the amount of the payment is determined by formula (rather than, for example, by the amount of money accumulated in an account). Thus, an arrangement whereby every employee reaching age 65 receives a pension of $100 a month for each year of service with the employer is a pension plan, as is an arrangement that provides a pension of 50% of final salary. In other words, we are dealing with "defined-benefit" plans and shall have little to do with "money-purchase" or "defined-contribution" plans, which provide for the deposit of a certain percentage of pay into an account for each employee each year, and where benefits are ultimately based on the balance in the account.

Pension plans are often festooned with ancillary benefits: survivor's pensions, lump-sum death benefits, disability benefits, and the like—but we shall not concern ourselves with these for the moment. Whatever else it may do, a pension plan always provides a life annuity to each participant who retires from employment (and who has satisfied the age and service requirements of the plan).

An employer who sets up a pension plan thereby commits the company to the expenditure of sums of money in future years. If the employer decides simply to pay the pensions as they fall due—*i.e.*, uses the "pay-as-you-go" cost method—then over the next 40 or 50 years the annual expenditures will show a pattern something like that of Figure 2.1.1. For the first several years (region A of the curve) the expenditures will rise only moderately, because at first few will be receiving pensions. During this initial period the employer enjoys the most pleasant years of pension sponsorship: The company gets credit from the employees for having established a fine, new plan, but the plan consists primarily of promises at this stage, so the cash outlay is small.

Typical Pattern of Benefit Payments from a Pension Plan

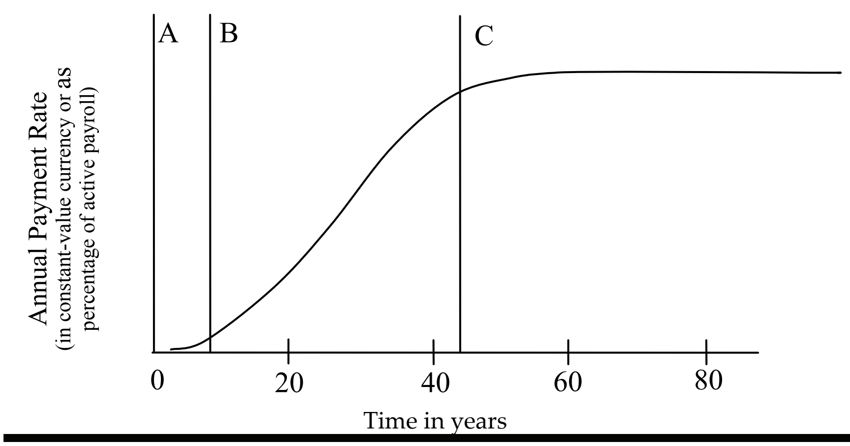

Figure 2.1.1

But when the passage of time brings the employee to region B of the curve, the plan starts to pinch. Employees keep retiring on bigger and bigger pensions, and those who are already retired are not dying at a sufficiently rapid rate to offset the new additions to the pension rolls. During this stage of pension–plan growth, annual outlays begin to rise steeply—outstripping the rise in covered payroll even if the employer stays in business at normal employment levels. This stage of growth is the most trying period for a pay-as-you-go pension plan, and is in fact the present stage of many governmental plans—U. S. Social Security in particular. During this stage the plan has not changed, there is no improvement in the benefit formula, but the cost is nevertheless snowballing, so the employer must pay ever increasing sums of money each year with no plan improvement to show for them.

At last, if the plan has survived the passage through region B, the cost begins to stabilize, more or less, as shown in region C of the curve of Figure 2.1.1. The danger is that the cost will stabilize at a level too high for the employer to bear. A pay-as-you-go pension plan depends on a thriving group of active employees to generate, either through taxes or through profit-making endeavor, sufficient revenues to pay the pensions. To make matters worse, the period of stability is sometimes followed by a period of decline in the fortunes of the plan sponsor, leading to a decline in the number of active employees, and in this case there is an even worse period following region C: The cost, measured against the number of active employees or total active payroll, explodes upward and the plan faces collapse.

The painful lesson that has been learned over and over again in the last century by various types of employer—first private employers and later public employers—is that the cost of a pension plan must be recognized during the *working lifetimes* of the employees who are ultimately going to receive pensions, preferably by actually setting aside funds sufficient to provide completely for each employee's life annuity at the time of retirement. When pension plans are funded in this manner, the safety of pensions being paid to those already retired is assured, and cannot be jeopardized by fluctuations in employment levels or by the financial collapse of the employer itself. So well has this lesson been learned that in the United States and Canada, and indeed in almost every modern industrialized country, it is generally not legal for a private employer to establish a pension plan that is not properly funded, or at least properly recognized on the books. Public plans, on the other hand, are sometimes still funded on a pay-as-you-go basis, but hardly anybody believes this is a sound state of affairs. (It is a political problem of huge proportions, however, when your predecessors have enjoyed the high praise and low cost of region A, while bequeathing to you the leapfrogging costs and no glory of region B, and to know that you have the unpalatable choice of either (a) raising your contributions right away to a level intermediate between your point on the curve and the ultimate level of region C, or (b) simply letting things ride along as they are and shipping the problem downstream to your successors.)

This is where actuaries come into the picture, because the outflow of benefit payouts over time is subject to life and other contingencies, and because the inflow of contributions is occurring at a different time and in a different pattern, so that the time value of money is an important consideration. The actuary can, by making assumptions about rates of return on the pension fund, ages at retirement, rates of turnover and mortality, *etc.*, assign to each fiscal year a portion of the present value of future benefit payments in such a way as generally to accrue costs over the working lifetimes of employees. Any scheme for making such an assignment of costs is called an *actuarial cost method*—which we shall henceforth refer to simply as a "cost method." There are many such cost methods in common use, each with a different philosophical foundation. We shall examine the most popular of these in the next pages—starting always from the philosophical underpinnings of each, and proceeding to a general description.

The application of a cost method to a particular plan in order to compute its cost is called an *actuarial valuation*. The same term applies to the process of determining the liabilities of an insurance company, but a pension-plan valuation differs from an insurance-company valuation in many ways—some obvious and others quite subtle—which will reveal themselves as our discussion unfolds. The pension valuation may involve the computation of "liabilities" and the valuation of assets, but its primary purpose is to determine annual cost.

Equation Section (Next)2.2 Unit Credit

Assuming that each employee is entitled to retire at age y with an annual pension (payable monthly) equal to $B(y)$, a properly funded plan should have accumulated for each employee when he reaches age y an amount sufficient to fund his pension, *i.e.*, an amount equal to $B(y)\ddot{a}_y^{(12)}$. This requirement is the first logical premise of the *unit credit* cost method (as well as a number of other methods, as we shall see).

Now the benefit $B(y)$ does not arise suddenly at age y, but is earned or "accrued" in a more or less continuous fashion during the employee's active years of service. Thus, when the employee is hired, say at age w, his accrued benefit $B(w)$ is exactly zero; at retirement age y it is equal to its ultimate value $B(y)$; and at any point in-between, say at age x, it has some intermediate value $B(x)$, which we call his *accrued benefit*.

At any age x, earlier than y, the present value of employee j's accrued benefit is equal to $B^j(x)\,{}_{y-x}p_x v^{y-x}\ddot{a}_y^{(12)}$. Note that the factor ${}_{y-x}p_x$ is computed using a table of q_x's which represents probabilities of termination of employment before age y from all causes—not just from mortality, but also resignation, discharge, disablement, *etc.* This table of q's is called a *service table*—a term parallel to but more general than *mortality table.*

So, if we had assets on hand at all times equal to $\sum_{\mathbf{A_t}} B^j(x)\,{}_{y-x}p_x v^{y-x}\ddot{a}_y^{(12)}$ [1]then no matter what was the distribution of ages among the group $\mathbf{A_t}$ of active employees at time t, we should be assured of having sufficient funds to be able to withdraw $B^j(y)\ddot{a}_y^{(12)}$ as each employee reached age y—even if all employees were the same age and all retired at once. (Of course, we might not actually withdraw money to purchase an annuity, but the philosophy is the same no matter what medium of funding is used. It will make our discussion clearer if we assume for the moment that retirees are removed from both the asset and liability columns of our pension plan. We shall put them back in later on.)

This observation is the source of the second premise of the unit credit cost method, which distinguishes it from all others: The ideal fund balance, or desired amount of assets, on hand at any given time t is equal to $\sum_{\mathbf{A_t}} B^j(x)\frac{D_y}{D_x}\ddot{a}_y^{(12)}$ (remember we are ignoring retirees).[2] This ideal fund balance is called the *accrued liability:*

$$(\text{Accrued liability})_{\text{t}} = AL_t = \sum_{\mathbf{A_t}} B^j(x)\frac{D_y}{D_x}\ddot{a}_y^{(12)} \qquad (2.2.1)$$

In other words, under the unit credit cost method, the accrued liability is defined as the *present value of accrued benefits.* This definition distinguishes it from all other cost methods, and carries with it, by implication, a complete definition of the pension cost that should be ascribed to any given year, as we shall now see.

Let us digress for a moment to remark on our peculiar use of the word "liability" to denote a desired level of assets. This oddity, which has caused no end of confusion among accountants,

---

[1] The ages x and y should also carry the superscript j, but we omit the superscript to reduce clutter.

[2] The definition of D_x may be found in the Index to Principal Notation.

arises from life-insurance terminology. In ordinary financial accounting, a business records each transaction twice—once on each side of the balance sheet—so its "liabilities" are, roughly speaking, the sum of amounts actually owed to someone else. In life-insurance accounting, by contrast, premiums received are not recorded on both sides of the balance sheet, but only as assets. To a life-insurance company, a "liability" is an actuarially determined amount that has first claim on the invested assets of the company. It is not, strictly speaking, an amount owed to anyone—although it will be if the reserve basis proves true—it is the amount of *assets* to be set aside for whatever the actual claims turn out to be. In the same way, the accrued liability of a pension plan represents a claim on plan assets.

From year to year the accrued liability changes, not only because the ages of the active participants increase, but also because the composition of the active group itself changes. To keep things simple, we shall assume that there are no new entrants into the plan; we shall put them in their own separate pension plan for the moment, and recall them later when we have need of them. Then the active group can never grow, but can only shrink, during the year. Denote by $\mathbf{T}$ the set of all employees who terminate employment between times t and $t+1$, and by $\mathbf{R}$ the set of employees who reach retirement age y during the year; then we can write

$$\mathbf{A}_{t+1} = \mathbf{A}_t - \mathbf{T} - \mathbf{R}. \tag{2.2.2}$$

We now construct the following purely algebraic argument to show the relationship between the accrued liability at time t and the accrued liability at time $t+1$ (using the results of exercise 2.2.1):

$$
\begin{aligned}
\textbf{(Accrued liability)}_{\mathbf{t+1}} \;&\equiv\; AL_{t+1} \\
&= \sum_{\mathbf{A_{t+1}}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} \\
&= \sum_{\mathbf{A_t}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} - \sum_{\mathbf{T+R}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} \\
&= \sum_{\mathbf{A_t}} B^j(x+1)\left[\frac{D_y}{D_x}(1+i) + q_x\frac{D_y}{D_{x+1}}\right]\ddot{a}_y^{(12)} \\
&\qquad - \sum_{\mathbf{T+R}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} \\
&= \sum_{\mathbf{A_t}}\left[B^j(x) + \Delta B^j\right]\frac{D_y}{D_x}\ddot{a}_y^{(12)}(1+i) \\
&\qquad + \sum_{\mathbf{A_t}} q_x B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} \\
&\qquad - \sum_{\mathbf{T+R}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)},
\end{aligned}
$$

where ΔB^j is the increase in j's accrued benefit during the year. This means that

$$
\begin{aligned}
AL_{t+1} \;=\; &\left[AL_t + \sum_{\mathbf{A_t}} \Delta B^j \frac{D_y}{D_x}\ddot{a}_y^{(12)}\right](1+i) \\
&-\left[\sum_{\mathbf{T}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)} - \sum_{\mathbf{A_t}} q_x B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)}\right] \qquad (2.2.3) \\
&- \sum_{\mathbf{R}} B^j(x+1)\frac{D_y}{D_{x+1}}\ddot{a}_y^{(12)}.
\end{aligned}
$$

Now look at the second bracketed term of (2.2.3). If actual experience during the year were in accord with assumed experience, this term would be zero. That is to say, the expected release of liability on account of termination of employment before age y

Referring still to Figure 6.3.1: At younger ages the shape of the curve is governed more by the employees' financial ability to retire than by physical inability to continue working. For example, if an employee had accumulated five or ten years' pay in savings, his propensity to retire would be much higher than if he had accumulated less than a year's pay. The nature of the job would also influence "retirement," which for some occupations (military service, for example) is just another word for change of occupation. Nevertheless, for any particular employee group we can imagine a function $\mu_x^{(r)}$ in the pristine setting where there are no retirement benefits from employment.

If we now introduce a retirement pension, such as Social Security old-age benefits, we in effect inject instant financial resources into an employee's life at whatever age the pension becomes payable. Suppose, for example, that a governmental old-age pension is payable to everyone in our group at age 65. Then $\mu_x^{(r)}$ would change and look something more like the curve in Figure 6.3.2.

Force of Retirement with Simple Plan

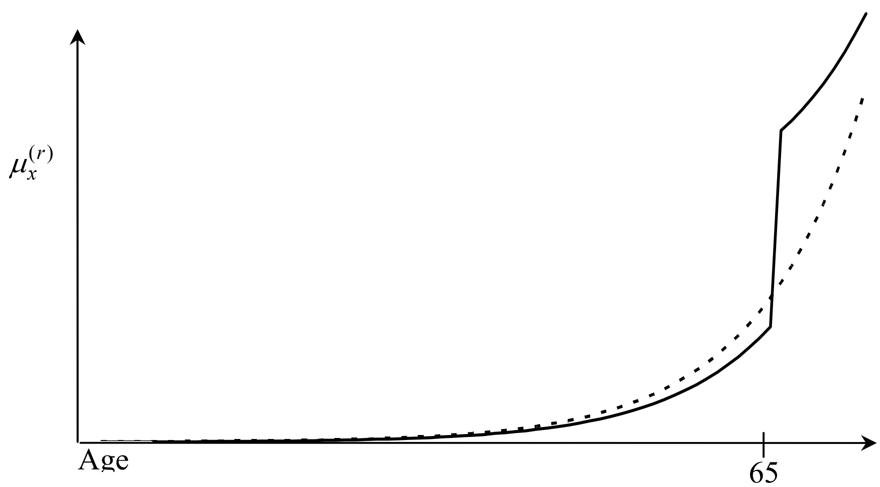

Figure 6.3.2

Close to age 65 the force of retirement is depressed below its "natural" value because employees tend to delay retirement until the government pension starts. Then at age 65 a sudden jump to a higher level takes place, and the force of retirement stays higher because the government pension is still available. Another way of looking at it is to imagine a whole family of natural $\mu_x^{(r)}$ curves similar to that of Figure 6.3.1—one for each level of financial resources: a group of relatively well-off employees would presumably fall on a higher curve than a group of poor workers. The hypothetical government pension simply shifts a particular group of employees into a higher resource category—like a quantum jump to a new energy level in atomic physics. The size of the vertical jump at age 65 would depend on the relative financial impact of the government pension: If it were low in relation to the worker's pay the jump would be smaller than if the pension were more generous.

When we now superimpose a private pension plan on the natural inclination to retire, as modified by the presence of the governmental system, we may introduce a number of such discontinuous jumps in the $\mu_x^{(r)}$ curve, corresponding to the various ages at which pensions become payable—for example, age 55 for early retirement under the private plan, age 62 for reduced social security, age 65 for full pension and full social security, and so forth. Also, the slope of the curve is modified in between each of these discontinuities because of the tendency of participants to wait for the next eligibility age and because the pensions add to the employee's wealth more rapidly toward the end.

A third variable affecting $\mu_x^{(r)}$, besides age and wealth, is direct pressure, either from the employer's mandatory retirement rules or a general peer pressure due to social custom (if your friends are all retiring at age 65 you tend to want to do likewise). Thus, in the real world the shape of $\mu_x^{(r)}$ tends to look more like Figure 6.3.3.

Force of Retirement with Complex Plan

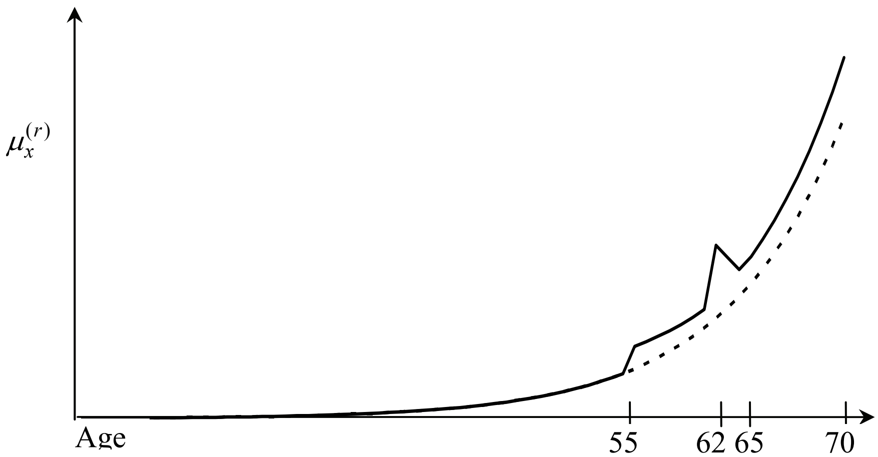

Figure 6.3.3

It is not necessary that $\mu_x^{(r)}$ increase between each of the discontinuities—it may get such strong "shocks" from benefit eligibility and the delaying effect that it actually decreases between, say ages 62 and 65.

If we were to continue this line of reasoning we would have to decide how much additional pressure to retire is added when each higher level of retirement benefits becomes available, and thereby be able to estimate the height of each jump in the curve. But such an approach would be of academic interest only, because the resulting probabilities would likely be no more accurate than a "seat-of-the-pants" estimate, by the actuary or the employer, of the probabilities of retirement themselves. Obviously, these probabilities vary with age and service and do not usually form a monotonically increasing sequence. They are also difficult to verify by actual experience, because peoples' inclination to retire has something to do with their faith in the future, and thus rests in part on the general condition of the economy.

In many situations, especially with plans whose early-retirement benefits are similar in value to the accrued liabilities, you will not gain accuracy by using a table of retirement probabilities instead of a single assumed retirement age. Assuming a single retirement age is tantamount to assuming a force of mortality like that illustrated in Figure 6.3.4.

Force of Retirement when Single Retirement Age Assumed

Figure 6.3.4

Note that under this assumption $\mu_x^{(r)}$ is zero everywhere except just before the assumed retirement age y, where it approaches infinity, thus ensuring that $q_{y-1}^{(r)} = 1$. In situations where you cannot justify assuming a single retirement age, such as where early retirement benefits are "subsidized," you have no choice but to assume retirement probabilities based on age, service, benefit level, and whatever other parameters may be important.

Whatever way you approach the probability of retirement, through use of a single age or a table of probabilities, the retirement decrement will close out the service table, in the sense that at some point $l_x = 0$.

Table 7.3.1

Comparison of Methods

	Individual Level Premium	Individual Aggregate	Modified Aggregate	Aggregate
Normal Cost	$\sum_{A_t} \frac{PVFB_t^j - AL_t^j}{PVFS_t^j} S_t^j$	$\sum_{A_t} \frac{PVFB_t^j - F_t^j}{PVFS_t^j} S_t^j$	$\frac{PVFB_t - F_t}{PVFNC_t} NC_t$	$\frac{PVFB_t - F_t}{PVFS_t} S_t$
Accrued Liability	$\begin{cases} 0 & t=0 \\ \sum_{A_t}(AL_{t-1}^j + NC_{t-1}^j)\frac{D_{x+t-1}}{D_{x+t}} & t>0 \end{cases}$	$\sum_{\mathbf{A_t}} F_t^j$; allocation Arbitrary	F_t	F_t
Unfunded	$AL_t - F_t$	**0**	**0**	**0**
Total Cost	$NC_t + \frac{UAL_t}{\ddot{a}_{\overline{n}\rceil}}$; *n* arbitrary	**Same as Normal cost**	**Same as normal cost**	**Same as normal cost**

Note: The symbols NC_t, AL_t, *etc.*, refer to the ILP method throughout the table.

Exercises

7.3.1 Show that under the I/A cost method

$$F_t^j = AL_t^j - \lambda_t^j PVFNC_t^j. \tag{7.3.22}$$

7.3.2 In this section we suggested two possible criteria for distinguishing "individual" cost methods from "aggregate" ones. Which of the following criteria are sufficient to distinguish an individual cost method from an aggregate one?

(a) (Under an individual cost method) j's accrued liability can never be negative.

(b) j's normal cost would be unchanged if no other employees were present.

(c) j's total pension cost (normal cost plus amortization of current unfunded) is entirely covered by his own normal costs before he retires, regardless of actuarial gains and losses.

(d) The accrued liability is not equal to current assets, except by accident.

(e) The actuarial gain can be determined without computing the current normal cost.

Try each of these criteria on all the cost methods so far studied (you may use Tables 7.1.1 and 7.3.1, and exercise 7.1.4, to refresh your memory). Can we classify cost methods unambiguously into the categories "individual" and "aggregate?"

7.3.3 Prove that assets are implicitly allocated to employees under the M/A method by Equation (7.3.21). Can any employee have *negative* assets?

7.3.4 Derive the following recursion relation for the loading factor under the M/A cost method:

$$\lambda_t = \frac{\lambda_{t-1}\left(PVFNC_{t-1} - NC_{t-1}\right)(1+i) - G_t}{PVFNC_t}. \tag{7.3.23}$$

7.3.5 (a) Throughout this section we have used the term "gain" to refer to G_t, the gain under the ILP method. Now that we have separate definitions of accrued liability and normal cost under the two new cost methods (see Table 7.3.1), how would you define the gain under the I/A method? Under the M/A method? [Hint: remember that "gain" means an unexpected decrease in the unfunded accrued liability, and take a look at Equation (7.3.17).]

(b) Derive an expression for the total I/A normal cost at time t in terms of its value at $t-1$ and your expression for the I/A gain from part (a). Does your normal cost remain a level percentage of salary if actual experience equals expected?

(c) Repeat part (b) for the M/A cost method.

7.3.6 Consider a plan covering just two employees, employee A and employee B. The plan provides a pension of one-half the final year's salary and is established 1/1/2008. On 1/1/2009 the plan still covers only A and B, but on 1/1/2010 you find that employee B has quit and been replaced by a new employee, C. The pertinent data on the three employees are as follows:

Name	A	B	C
Date of birth	1/1/1958	1/1/1968	1/1/1978
S_0^j	\$50,000	\$20,000	
S_1^j	60,000	25,000	
S_2^j	70,000		\$22,000

The actuarial assumptions are as follows:

Interest	5%
Service table	None
Salary increases	None
Retirement age	65
$\ddot{a}_{65}^{(12)}$	10.0

The employer actually deposits each contribution (with one year's assumed interest into the fund at the end of each year. The fund earns 10% the second year.

(a) Compute the costs of the plan as of 1/1/2008, 1/1/2009, and 1/1/2010 using the ILP method with each year's gain amortized over a new 15-year period.
[Answers: \$13,029, \$15,994, \$16,952]

(b) Compute the costs for the three years under the individual aggregate method using the gain-allocation method of Equation (7.3.13).
[Answers: \$13,029, \$15,994, \$16,926.]

(c) Compute the three costs under the modified aggregate method. [Answers: \$13,029, \$15,994, \$16,926.]