

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Pensions, Labor, and Individual Choice

Volume Author/Editor: David A. Wise, ed.

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90293-5

Volume URL: <http://www.nber.org/books/wise85-1>

Publication Date: 1985

Chapter Title: The Riskiness of Private Pensions

Chapter Author: Jerry R. Green

Chapter URL: <http://www.nber.org/chapters/c7139>

Chapter pages in book: (p. 357 - 378)

12 The Riskiness of Private Pensions

Jerry R. Green

12.1 Introduction

Private pensions represent the contract between workers and their employers. The future benefits of the pensions are an asset for the workers who hold them. The contractual obligations of the plan to the worker are quite complex. Benefits are due in a variety of circumstances, in each of which a complicated formula may be employed to determine the amount to be paid. Plans differ both in the description in these circumstances and in the nature of these formulas.

It is important to recognize that the pension does not represent a claim to a perfectly certain stream of resources. Plans with differing provisions will differ both in the mean payoff and in its variability. Moreover, when computing the mean and the variability of pension benefits the age of the worker and his length of job tenure should be considered. Different plans may favor different sets of workers.

The object of this paper is to discover how the alternative forms of the pension contract affect these means and variabilities for workers of different circumstances. Specifically, I aim to provide some information relevant to the following questions. How should workers evaluate respective benefit packages offered to them by different plans? How does this assessment depend on the workers' expectations of economic risks such as fluctuations in wages, interest rates, and inflation? How does it depend on their beliefs about longevity, disability, retirement, and their propensity toward mobility in employment? Will the evaluation of different plans differ markedly for workers of different ages and job tenure?

Jerry R. Green is professor of economics at Harvard University and a research associate at the National Bureau of Economic Research.

I gratefully acknowledge excellent research assistance by Arturo Estrella and helpful comments from Alan Auerbach, Peter Diamond, and Herman Leonard.

It may be useful at the outset to reflect on the sources of risk faced by a worker enrolled in a pension plan. Some risks, such as fluctuations in the price level or interest rate, are common to all workers living at the same time. Others, like the future wage of a worker on which a pension may be based, are correlated across workers covered in the same plan. They are also positively correlated for workers in different plans or different industries. Finally, individual-specific risks such as death and disability are unlikely to be strongly correlated even among workers in the same plan.

In addition, we must consider risk factors that are not exogenous to the individual. For a worker with vested benefits, his decision to leave the firm would cause a reduction in the benefits he would receive from his pension plan on retirement. Moreover, for those plans with an early retirement provision, such a decision would change the value of accrued benefits. Workers who are not fully vested face even more of a loss when they voluntarily terminate employment.

The complexity of pension plans and the multifaceted nature of the risks mentioned above make it impossible to conduct a purely theoretical analysis of the questions in which we are interested. Therefore, the methodology of this paper will be to use simulation analysis. I compare four different forms of a pension contract. I simulate the economically relevant events during the worker's lifetime, both macroeconomic events and the individual-specific events mentioned above, as they evolve during his lifetime. To the extent possible I have used current actuarial projections.

The analysis in this paper has importance for studying the impact of private pensions in the national economy. Most studies to date have regarded the pension asset as a claim to a perfectly certain stream of consumption for the worker. It is clear that this is far from the case. In studying saving behavior, for example, it is important to distinguish between individuals' holdings of financial assets in their own name and the amount of their private pension wealth. The two are far from perfect substitutes.

Similar considerations apply for the issuers of pensions. Firms have an implicit liability to their pension holders. The riskiness of this liability has not as yet been treated. Because pensions cover a large number of workers, the law of large numbers mitigates those components of the risk that are due to individual-specific events. However, macroeconomic risks do affect the realized present value of aggregate plan liabilities. Therefore, even if pensions were recognized by both sides of the "pension market" as a form of contractual arrangement that should be incorporated into their respective wealth calculations, the net effect on national savings might very well be nonzero because the relevant risks cannot be perfectly insured. The certainty equivalent value of a pension for workers in the plan is probably much less than the value of the liability that it represents to the stockholder in the firm. Therefore, distributional considerations aside, pensions decrease the value of other forms of saving by workers by less

than they increase the required saving by stockholders necessary to offset their indirect pension liabilities.

By providing information about the risks inherent in pension contracts this paper will be a useful and important ingredient in the analysis of the role of private pensions in the economy. In addition, by comparing different plans with respect to their riskiness, insight might be gained into the problem of designing new plans and into the selection of benefit changes in existing plans that would be most beneficial to their members.

12.2 Description of Plans to Be Compared

I compare four hypothetical forms of pension plans. These plans differ both in the way that they compute benefit levels for retired workers and in the circumstances other than retirement under which benefits are paid. There are three types of these payments: early retirement, disability, and death benefits. Some of the plans provide benefits to a surviving spouse in the event of preretirement death. None provides postretirement death benefits.

The four plans studied also differ slightly in the vesting formulas used. Basically, they provide for 10-year vesting, but two of them gradually increase the vesting level from 50% at the tenth year to 100% at the fifteenth year of service.

The most important difference among the plans is in the way the benefits are computed.

In Plan 0, a “conventional” plan, benefits are proportional to the average of the highest 10 annual compensations. The proportion is increasing linearly with years of service. In Plan 1, a “pattern” plan, the benefit is a fixed dollar sum per year of service. Plans 2 and 3 are more complex. Plan 2 is a “career average” plan. The benefit is a fixed proportion of each year’s compensation. This plan is favorable to workers with long tenure in the firm. Plan 3, a “final salary basis” plan, is a mixture of a proportion of the individual’s highest salary plus a bonus for workers with very long tenure. In addition, Plan 3 uses the worker’s best four years as a base for the computations, whereas Plan 2 bases benefits on the compensation received in all years.

Table 12.1 summarizes these provisions.

12.3 Description of the Simulation of Benefits Received

12.3.1 General Description

Two sets of simulations were conducted. The first takes the plans precisely as specified above. The second suppresses all pension benefits except the retirement benefit. This enables us to focus more clearly on the

Table 12.1 Plans

Plan	Vesting	Early Retirement	Disability	Death Benefit (Preretirement)	Benefit
0 Conventional	10 years	—	—	—	1½% per year of service times average of highest 10 annual compensation
1 Pattern	10 years	a)Age 55 Pension minus 3% for each year under 65 (or if less, by 3% for 85 – age + service) b)Age 62 Full pension if 10 years' service	10 years' service Full pension	Age 55 or more Spouse receives 50% of early retirement benefit for life	\$144 per year for each year of service
2 Career average	10 years: 50% + 10% for each additional year. 5 years if age + service ≥ 45: 50% + 10% for each additional year. Age 60: 100%	Age 55 and 10 years' service. Reduce pension by 2.2% per year for first 3 years below 65; 4.8% per year thereafter	Age 55, 10 years' service. Reduce pension by 2.2% per year for first 3 years below 65; 4.8% per year thereafter	None	½% of 1973 compensation times years of service prior to 1975 + ¾% of each subsequent year's compensation
3 Final salary	10 years: 50% + 10% for each additional year. 5 years if age + service ≥ 45: 50% + 10% for each additional year. Age 60: 100%	Age 50 and 15 years' service. Pension reduced by 6½% for each year age if less than 65	None	Age 50 and 15 years Spouse receives 50% of early benefit	55% of annual compensation in highest 4 consecutive years + ½% extra for each year of service over 30, up to maximum of 65% of average annual compensation in 4 highest consecutive years

differences in the retirement benefit formulas. I shall now describe how the common components of these two simulations were set up. These include the macroeconomic risks, the worker's wage process, and the individual risks to which each worker is exposed.

The simulation of benefits is based on the presumption that the structure of the plan will remain fixed. The plans' basic provisions for computing benefits determine their risk characteristics. In order to calculate benefits for Plans 0, 2, and 3, we need to know the worker's wage history. Plan 1 does not require a wage history because it is a flat benefit per year of service. The wage process, therefore, is the first essential piece of information that must be simulated.

Although we assume that the structure of plans is fixed, we do not assume that the precise provisions are immutable. Quite to the contrary, we presume that the level of benefits specified in the pattern plan, Plan 1, will be updated so as to reflect the rise in real wages of the average worker, over time. It is important to note, therefore, that it is the average worker's wage that affects Plan 1 payoffs, whereas it is the individual worker's wage that is an input into the other calculations.

The risk characteristics of Plan 1 are different from the other plans for this reason. Workers in Plan 1 are at risk for fluctuations in the difference between the adjusted dollar benefit level and actual inflation. But they are not at risk for the difference between their own personal compensations and the average over all workers in the plan.

The demographic factors are the second major ingredient in the simulations: job leaving, disability, death, and death of a spouse where this is relevant.

My procedure is as follows. The age and years of experience of an individual are fixed. For example, assume that the individual is now 30 years old and has already been in the plan for five years at the time when the simulation is performed. It is assumed that the worker receives compensation in the current year. Thereafter, annually, I simulate the wage of all workers, his wage, and his demographic status. If he or his surviving spouse is due benefits under any plan, I keep track of them. Thus, at the end of each individual's lifetime I have a record of the stream of benefits received and the circumstances under which they were paid. From this series one can compute, retrospectively, its present value and its present values conditional on the occurrence of various events.

For every age/experience situation, and each plan, I run a sample of 900 individuals through the simulated lifetime described above. Experimentation has indicated that this size is enough to eliminate most of the sampling variance and to provide fairly accurate evidence on the risks due to the random factors driving the simulation.

In the next two subsections I provide the details of the simulation methods for the wage rates and the demographic factors, respectively.

12.3.2 Simulation of Wage Process

The initial wage in all cases was taken to be \$20,000. Because I am primarily interested in the risks faced by workers relative to their total wages or total compensations, one can regard this merely as a normalization.

The growth rate of average wages from t to $t + 1$ is denoted \bar{w}_t . For the individual worker it is w_t . Suppose that

$$w_t = \bar{w}_t + \eta_t,$$

where η_t is a normal random variable, independently and identically distributed over time. The parameter η is assumed to have a mean of zero and a standard deviation of .005.

As the individual ages, shocks via η_t are cumulative. It is the idiosyncratic component of rate of growth of his wages that evolves independently.

The stochastic structure of \bar{w}_t is autoregressive: $\bar{w}_t = \rho \bar{w}_{t-1} + u_t$.

The parameter ρ is taken to be .9 and the mean and standard deviation of u_t are assumed to be .001 and .002179, respectively. Thus the mean of \bar{w}_t is .01, as is the mean of w_t .

For Plans 0, 2, and 3, the resulting individual wage for a worker with s years of service in the future is

$$w_s = 20,000 \left[\prod_{t=0}^s (1 + w_t) \right], s = 1, 2, \dots$$

The series $w_0, w_1, \dots, w_s, \dots$ forms the basis of compensation and the benefit package.

In Plan 1, a different computation for pension benefits is used, based on average wages rather than on the personal wage of the individual in question,

$$\bar{w}_s = 20,000 \left[\prod_{t=0}^s (1 + \bar{w}_t) \right], s = 1, 2, \dots$$

and the annual retirement benefit is simply

$$\$144 \times \frac{\bar{w}_T}{20,000} \times T,$$

where T is the number of years of service at retirement.

Because of the positive drift in real wages built into this process, virtually all workers will have a higher real wage on retirement than they do at the date when the simulation is undertaken. However, because of the cumulative effect of η_t , which represents promotions and other random variations due to the profitability or lack thereof in the firm or industry in question which may affect the wages of its workers, there will be a substantial variation across workers, and increasingly so as time passes.

I initiate the simulation for various ages of workers, in each case trying to assess the remaining risk in the plan benefits and their correlation with

the worker's wage. The parameter specifications above determine the distribution of the worker's terminal wage, at age 65, and how its variance breaks down into the part due to aggregate real wage fluctuation and the idiosyncratic component. This breakdown is shown in the unnumbered table below.

Age	EW_{65} 20,000	Standard Deviation of EW_{65} 20,000	Fraction of Variance Due to Idiosyncratic Risk
25	1.43	.152	.13
30	1.45	.168	.11
35	1.35	.161	.08
40	1.31	.126	.11
45	1.21	.093	.13
50	1.15	.073	.11
55	1.09	.055	.23
60	1.05	.028	.37

12.3.3 Simulation of Events in Workers' Lifetimes

The possible states in which a worker or his beneficiary can be found within my simulation are shown in figure 12.1 below.

It can be seen that the passage through this flowchart is governed by the probabilities of transition from working to the three other states. In the first set of simulations, I assume that these are a function of age only, not

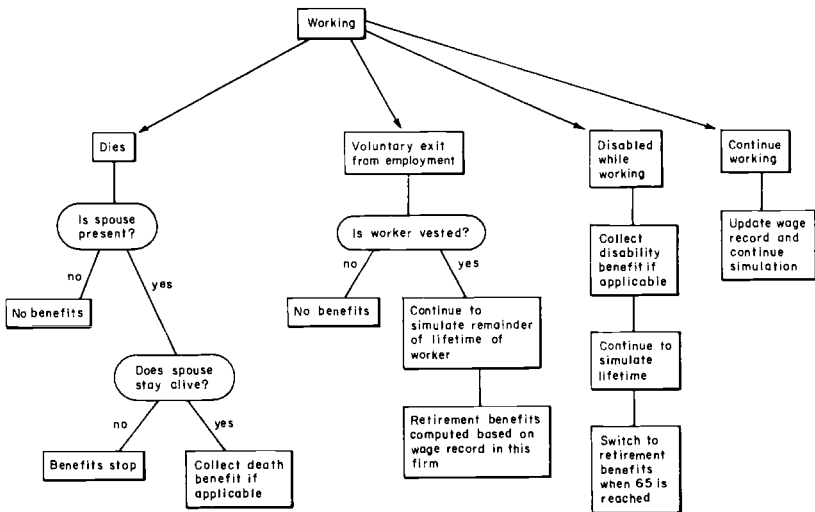


Fig. 12.1 One step in simulation of a worker's pension benefits.

of experience. In the second set of simulations, the probability of withdrawal from employment is allowed to depend on both age and experience. Table 12.2 gives the probabilities used.

For the simulation of spouses' lifetimes, I used the age-specific mortality for men shifted by eight years. Three years represents the average age differential of husband and wife, and the remaining five years is an attempt to capture the different mortality probabilities for men and women. It approximates standard actuarial practice.

Table 12.3 presents the assumed mortality table, up to age 110, where the simulation of lifetimes was always terminated.

The specific features of this flowchart embody the characteristics of the benefit packages studied and some other special features owing to the limitations on our actuarial knowledge.

1. No early retirement is provided for in this flowchart, even though Plans 1, 2, and 3 have early retirement provisions. The reason is that I do not have estimates of early retirement probabilities. Many "disabilities" are probably early retirements of a certain kind. To incorporate early retirements separately would require some reduction in the assumed disability probabilities, and the extent of this is beyond my knowledge.

Moreover, the early retirement benefits provided by these plans are arranged to approximate an actuarially fair payment of the accrued benefit. For example, if a worker age 55 retires in Plan 1, he receives 70% of the pension he would have received at age 62. If he were simply to withdraw, he would get these benefits in full beginning at age 62. Depending on his mortality, the value of that stream could be higher or lower than the early retirement benefits.

Finally, I suspect there is a substantial amount of adverse selection in the early retirement decision. Workers who know that they are in poor health or are unlikely to have long lifetimes for some other reason are more likely to take this option. It would therefore be a mistake to use the same mortality tables for this self-selected population as for the workers as a whole. Because the magnitude of this effect is unknown, it seemed simpler to leave these workers out entirely.

2. No plans pay spouse benefits except in the event of preretirement death. Therefore the presence or absence of a spouse and the subsequent lifetime of the spouse need only be simulated down that branch of the tree.

3. For workers who withdraw from the plan after they have vested benefits, no provisions of the plan are operative except for the retirement benefit at age 65. Therefore, for these workers I need not keep track of their subsequent employment or disability status. Only their lifetimes need to be simulated.

4. The disability benefits provided by Plans 1 and 2 are paid only to vested workers. Assume that the spouse of a disabled worker who subse-

Table 12.2 **Age-Specific Probabilities of Death, Withdrawal, and Disability as Assumed in the Simulation (Percentage of the Population at Risk in the Indicated Age Group)**

Age	Death	Withdrawal	Disability
20	.0044	24.30	.0263
21	.0463	22.44	.0266
22	.0488	20.70	.0270
23	.0513	19.07	.0271
24	.0538	17.56	.0274
25	.0569	16.15	.0275
26	.0602	14.85	.0278
27	.0636	13.64	.0281
28	.0679	12.53	.0279
29	.0717	11.51	.0283
30	.0763	10.58	.0382
31	.0817	9.73	.0382
32	.0872	8.95	.0381
33	.0940	8.26	.0382
34	.1006	7.63	.0386
35	.1080	7.08	.0386
36	.1164	6.57	.0487
37	.1261	6.13	.0586
38	.1360	5.74	.0680
39	.1475	5.40	.0781
40	.1593	5.11	.0873
41	.1748	4.86	.0971
42	.1945	4.65	.1177
43	.2208	4.47	.1365
44	.2509	4.32	.1566
45	.2852	4.20	.1762
46	.3240	4.09	.1952
47	.3680	4.01	.2156
48	.4131	3.93	.2442
49	.4640	3.87	.2751
50	.5165	3.80	.3034
51	.5757	3.74	.3339
52	.6341	3.68	.3722
53	.6971	3.60	.4108
54	.7650	3.52	.4493
55	.8508	0	.4971
56	.9238	0	.5362
57	.9997	0	.5966
58	1.0859	0	.6762
59	1.1868	0	.7968
60	1.3042	0	.9721
61	1.4337	0	1.2319
62	1.5748	0	1.5857
63	1.7329	0	2.0619
64	1.8932	0	2.6739

Source: H. E. Winklevoss, *Pension Mathematics, with Numerical Illustrations* (Homewood, Ill.: R. D. Irwin, 1977).

Table 12.3 1971 Male Group Annuity Mortality Rates

Age	Mortality	Age	Mortality	Age	Mortality	Age	Mortality	Age	Mortality
20	.00050	40	.00163	60	.01312	80	.08743	100	.32983
21	.00052	41	.00179	61	.01444	81	.09545	101	.35245
22	.00054	42	.00200	62	.01586	82	.10369	102	.37722
23	.00057	43	.00226	63	.01741	83	.11230	103	.40621
24	.00059	44	.00257	64	.01919	84	.12112	104	.44150
25	.00062	45	.00292	65	.02126	85	.13010	105	.48518
26	.00065	46	.00332	66	.02364	86	.13931	106	.53934
27	.00068	47	.00375	67	.02632	87	.14871	107	.60609
28	.00072	48	.00423	68	.02919	88	.15849	108	.68747
29	.00076	49	.00474	69	.03244	89	.16871	109	.78543
30	.00081	50	.00528	70	.03611	90	.17945	110	1.00000
31	.00086	51	.00587	71	.04001	91	.19049		
32	.00092	52	.00648	72	.04383	92	.20168		
33	.00098	53	.00713	73	.04749	93	.21299		
34	.00105	54	.00781	74	.05122	94	.22653		
35	.00112	55	.00852	75	.05529	95	.24116		
36	.00120	56	.00926	76	.06007	96	.25620		
37	.00129	57	.01004	77	.06592	97	.27248		
38	.00140	58	.01089	78	.07260	98	.29016		
39	.00151	59	.01192	79	.07969	99	.30912		

Source: Winklevoss, *Pension Mathematics* (see table 12.2 n.).

quently dies receives nothing, although the specification of Plan 1 is arguably vague on that point (does “preretirement” death imply that the worker is actively working at that point in time?).

The result of this simulation is that workers end up in one of the four branches of the tree depicted in figure 12.1. The probabilities of each of these four terminal states, for workers of differing initial ages are shown in table 12.4.

12.4 Results of Simulations: First Set

The first set of simulations treats all the benefits from the four plans exactly as specified in section 12.2. Table 12.5 shows the expected present values in each of the plans from the point of view of workers at various ages and various levels of experience. In all circumstances, Plan 3 is the most generous and Plan 0 is second.

Interesting differences appear between Plan 1, the pattern plan, and Plan 2, the career average plan. The former is much better at the older and more experienced end of the spectrum and the latter is superior for younger and less experienced workers. Both age and experience are important determinants of the relative benefits. The 45-year-olds with only five years of experience clearly prefer the career average plan, while those with 20

Table 12.4 Probability of a Worker's Ending the Simulation in Each State

Age	Death	Withdraw from Plan	Disabled	Retires with Benefits at 65
20	.010	.960	.006	.024
25	.026	.877	.018	.079
30	.050	.748	.036	.167
35	.075	.601	.055	.269
40	.096	.456	.073	.375
45	.113	.316	.087	.484
50	.121	.167	.098	.613
55	.118	0	.103	.779
60	.075	0	.080	.846

Note: In the second set of simulations, a more complex simulation of the decision to withdraw voluntarily from the plan is made. This allows years of experience in the plan as well as age to affect the withdrawal probability. The probabilities of death and disability remain exclusively age dependent. This change results in only minor shifts in the cumulative withdrawal probabilities, predictably being lower for the more experienced workers within each age cohort.

Table 12.5 Expected Present Values of Pension Benefits

Plan Type	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience				
Age 25	21666	8380	13456	32087
Age 40	40205	20264	33109	95554
5 years' experience				
Age 30	50826	20815	27071	73762
Age 45	57560	26435	34233	121497
10 years' experience				
Age 35	76947	34612	38475	114711
Age 50	75228	36365	38343	137765
20 years' experience				
Age 45	109853	51683	49751	138026
Age 60	89440	47704	39095	154902

Note: Flow of benefits over worker's and spouse's lifetime at 1% real discount rate.

years of experience have a preference for the pattern plan, risk considerations aside.

Table 12.6 shows the correlations between the present value of the wage stream remaining and the present value of benefits as specified in the plan. Quite naturally there is a substantial correlation in Plans 0, 2, and 3, where benefits are proportional to some sort of salary average, although the precise provisions vary from one to the other. But note also that Plan 1 displays a high correlation, indeed, the highest at the younger ends of the

Table 12.6 Correlations between Wages and Pension Benefits

	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience				
Age 25	.78	.80	.78	.77
Age 40	.66	.76	.71	.66
5 years' experience				
Age 30	.76	.76	.77	.67
Age 45	.67	.68	.58	.42
10 years' experience				
Age 35	.66	.62	.65	.47
Age 50	.55	.48	.45	.30
20 years' experience				
Age 45	.52	.46	.48	.38
Age 60	.37	.19	.31	.12

spectrum. This is due to the simulation of the real wage process as highly autoregressive. The present value of an individual's wages is substantially influenced by the general level of wages in the economy in which he lives. Those with favorable experiences, in my simulation, enjoy the benefits of a high wage as individuals, on average, as well as a high level of pension benefits because these are proportional to the average wage. This effect is attenuated only for the oldest workers—those at 60 years old with 20 years of experience. The effective fluctuation in their benefits is sharply curtailed because the level of benefits for any one individual is assumed not to be adjusted after age 65. Subsequent variations in workers' real wages do not induce any dispersions in the benefits received by a cohort of retired workers. Moreover, having already lived to 60, the largest portion of the fluctuations in wages is already behind them, and the wages remaining vary only slightly within the final five working years.

Table 12.7 shows the risks inherent in the various plans for the total compensation of workers. Because the plans differ in the level of benefits provided, some adjustment had to be made to equate their average returns before a meaningful comparison of risks could be attempted. The procedure I used is as follows. The benefit levels of Plans 1, 2, and 3 were factored up by an amount such that their average payoff to all workers is equal to that in Plan 0. The age/experience distribution of the work force is assumed to be uniform over the eight classes. Then, the coefficient of variation of the present value of wages, pension benefits, and their sum was computed for each plan and each age/experience category. This procedure is justified by the idea that the expected costs of the benefit package, as viewed by the firm, are the relevant aspect of the bargaining agreement. Such a bargain, whether explicit or implicit, is struck collectively, covering all workers in the firm whatever the age and experience they happen to have at the time.

Table 12.7 Ratios of Standard Deviations to Means: Wages, Pensions, and Total Compensation

	Wages*	Plan 0: Conventional		Plan 1: Pattern		Plan 2: Career Average		Plan 3: Final Salary	
		Pension	Total	Pension	Total	Pension	Total	Pension	Total
0 experience									
Age 25	1.185	2.477	1.266	2.256	1.246	2.447	1.308	2.132	1.236
Age 40	.587	1.278	.632	1.124	.649	1.114	.667	.945	.624
5 years' experience									
Age 30	.874	1.433	.918	1.354	.897	1.458	.938	1.168	.871
Age 45	.534	1.036	.577	.923	.566	.890	.559	.683	.498
10 years' experience									
Age 35	.736	1.000	.736	.854	.707	.996	.738	.740	.671
Age 50	.427	.794	.450	.625	.420	.714	.437	.593	.391
20 years' experience									
Age 45	.521	.812	.531	.675	.492	.781	.520	.656	.478
Age 60	.182	.604	.340	.518	.308	.592	.337	.508	.289

*Wages are the same in every plan simulation.

The results are quite unambiguous. Plan 3, the final salary plan, is the safest at all ages and experience levels. This is somewhat surprising at first glance because the level of benefits depends heavily on events at the end of the working life. It might seem that Plan 3 is considerably more risky than plans, such as the career average plan, that allow a wider dependence on such stochastic phenomena and presumably more pooling of risk over the workers' lifetimes. The reason for the result, as can be seen in table 12.2, is the low correlation between these benefits and lifetime wages. A worker who wants a safe lifetime earnings package should prefer a plan whose benefits are not overly sensitive to his own lifetime wages. This is precisely why Plan 3 does well.

One might ask why the pattern plan (Plan 1) is not best in this regard, since its benefits are fixed and independent of the worker's wage. Indeed, Plan 1 ranks as the second safest at every age/experience class except the 45-year-olds with five years of experience, where the difference is very small. The reason might be that the pattern plan is based on terminal average wages, whereas the so-called final salary plan is based on an average of four years' salary, which allows an intertemporal pooling of wage risk. In addition, however, other types of benefits are included in the plans, and Plans 1 and 3 differ in these regards. Plan 1 provides a disability benefit and Plan 3 does not, and there is a slight difference in their preretirement health benefits. To sort out the components of these risks more carefully we present a set of conditional calculations below.

Economic theory and common sense tell us that the risks inherent in any asset cannot be assessed without reference to other random factors with which they may be correlated. The calculations above are based on a presumption that the worker cannot offset any of his wage or pension risk by other market actions. Thus far we have neglected the fact that many of the fluctuations in pension benefits are in direct response to events in the worker's lifetime.

The termination of benefits on the death of a retired worker represents an annuity aspect of the pension contract that may well be preferable to its equivalent in terms of expected present value, paid in cash. Similarly, the need for money in the event of disability or preretirement death is not the same as if the individual were continuing to work within the same firm.

At the most detailed level of generality the utility function should be conditioned on the "state" realized by the worker, the state being a full description of all relevant events in his or his family's life. As an approximation of this, I present some calculations of the riskiness of various plans conditional on certain categories of lifetime experience. Specifically, I look at those subsets of my 900 simulated individuals in each age/experience category that reach the states of disability, voluntary withdrawal, preretirement death, or retirement. I make no attempt to distinguish individuals according to the age at which the first three of these events take

place. Needless to say, this approximation is a coarse one because the severity of the loss in future income, from the surviving family's point of view, is quite sensitive to their stage in the life cycle.

Tables 12.8, 12.9, 12.10, and 12.11 present these conditional calculations for four events. Individuals are grouped into disability, preretirement death, and withdrawal categories. In table 12.11 I group those who have withdrawn from the plan together with those who retire. These individuals are likely to have had similar experiences regarding health and lifetime work patterns. To this extent they are in the same state. But it

Table 12.8 Ratios of Standard Deviations to Means: Conditional on Preretirement Death

	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience:				
Age 25	.440	.469	.440	.495
Age 40	.357	.441	.512	.478
5 years' experience:				
Age 30	.496	.536	.496	.521
Age 45	.417	.442	.417	.479
10 years' experience:				
Age 35	.440	.486	.440	.507
Age 50	.462	.495	.462	.535
20 years' experience:				
Age 45	.487	.544	.487	.553
Age 60	.458	.433	.458	.477

Table 12.9 Ratios of Standard Deviation to Means: Total Compensation Conditional on Withdrawal

	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience:				
Age 25	1.082	1.074	1.135	1.091
Age 40	.698	.739	.769	.762
5 years' experience:				
Age 30	.833	.825	.856	.818
Age 45	.701	.710	.532	.491
10 years' experience:				
Age 35	.663	.629	.644	.583
Age 50	.418	.389	.362	.362
20 years' experience:				
Age 45	.488	.434	.402	.404
Age 60	—	—	—	—

Table 12.10 Ratios of Standard Deviations to Means: Total Compensation Conditional on Disability

	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience:				
Age 25	.202	.149	.233	.182
Age 40	.395	.389	.407	.385
5 years' experience:				
Age 30	.336	.294	.369	.315
Age 45	.435	.404	.426	.379
10 years' experience:				
Age 35	.381	.293	.419	.353
Age 50	.475	.378	.454	.433
20 years' experience:				
Age 45	.349	.317	.352	.317
Age 60	.382	.305	.299	.355

Table 12.11 Ratios of Standard Deviations to Means: Total Compensation Conditional on Retire or Withdraw

	Plan 0: Conventional	Plan 1: Pattern	Plan 2: Career Average	Plan 3: Final Salary
0 experience:				
Age 25	1.274	1.249	1.312	1.235
Age 40	.633	.643	.649	.608
5 years' experience:				
Age 30	.936	.912	.953	.882
Age 45	.555	.550	.525	.463
10 years' experience:				
Age 35	.731	.700	.725	.652
Age 50	.379	.267	.359	.319
20 years' experience:				
Age 45	.457	.428	.441	.409
Age 60	.246	.266	.253	.252

must be recognized that those younger workers who withdraw from the firm will have other earnings and presumably other pension benefits. Thus the dispersion in total compensation induced by their withdrawal is likely to be highly negatively correlated with other earnings, and to that extent the figures presented in the table are an overestimate of the risks inherent in the plan. The comparison across plans, however, should be relevant. And for older workers the "retire or withdraw" category consists almost entirely of retirements.

Table 12.11 indicates a considerable degree of similarity between these conditional results and the unconditional results of table 12.3. Plan 3 is

the safest for almost every age/experience category. Only the oldest group, 60-year-olds with 20 years of experience, find Plan 0 safer than Plan 3. And for this group the four plans are so similar that the discrepancy is probably within the sampling error of the simulation.

The other conditional tables 12.8, 12.9, and 12.10 display a wider divergence from the unconditional calculations shown in table 12.7. Preretirement death benefits are paid only in Plans 1 and 3. Therefore all of the variances in total compensation in Plans 0 and 2, conditional on preretirement death, arise from the distribution of wages and the dispersion in working intervals until the death occurs.

Disability benefits are most generous under Plan 1. Since disability results in a termination of wages, the variance of total compensation is reduced to the extent that benefits compensate for this loss in income. Thus Plan 3 does poorly because it incorporates no provision for disability. Plan 2 is safe for older, experienced workers because they receive fuller compensation, but it does poorly for those disabled early in their careers because their benefits are sharply curtailed.


12.5 Results of Simulations: Second Set

Taking the pension plans as written and simulating their ultimate benefits, as is done in section 12.4 has left us with some unresolved issues. The variance in total benefits depends very much on benefits other than those paid at retirement. In order to compare benefit formulas more directly, I shall redo the simulations, suppressing all benefits except those received at retirement.

The variability in payoff still depends on all the stochastic factors mentioned in sections 12.3 and 12.4. When these simulations were run, a more elaborate and realistic probability structure for withdrawal was incorporated. Voluntary withdrawal from employment depends on experience within the firm as well as on age.¹ This change produces slightly different frequencies in the terminal states reached by the 900 people whose lifetimes are simulated, but they are small relative to the sampling errors in this calculation.

This simulation allows a cleaner comparison of the four pension benefit formulas. I look first at the mean benefits. As in section 12.4, I adjust the level of payoffs so that the average mean payoff is equalized for a population equally distributed over the 10 age/experience groups. Then I compare the mean retirement benefit as it would be prospectively viewed by the representative worker in each group. The results are shown in figure 12.2.

1. I am grateful to Peter Diamond for references to the work of Yves Balcer that were very useful at this point.



Each point shows the benefits received at retirement, on average, in the indicated plan. Points marked ● are the benefits for the workers whose experience is 25 years less than their age. Points marked + are for workers with experience 40 years less than their age. The number indicates the plan. Benefit levels are adjusted so that the average overall ten age/experience groups is the same in each plan.

Fig. 12.2 Mean present value of benefits in the second set of simulations.

Younger workers find Plan 2, prospectively, the most favorable. Middle-aged, experienced workers rank Plan 0 highest. But their contemporaries with little experience rank it poorly. The final salary plan, Plan 3, does very badly for the younger workers and very well for the older ones.

It would be interesting to see whether the form of benefit calculation was systematically related to whether a well-defined subset of the workers had the predominant share of the negotiating power. Presumably this would vary across industries and across union versus nonunion firms. This must be left for future research.

The second use to which this simulation is put is a reassessment of the variability comparisons performed in section 12.4. The results are reported in tables 12.12 and 12.13. Because only retirement benefits are included, I report the conditional coefficient of dispersion of these benefits for retired persons only. These results in table 12.13 reveal the same general pattern as the comparable calculation in table 12.11. Plan 3, the final salary plan, is the safest throughout almost the entire age/experience spectrum. It is only for the 40-, 45-, and 50-year-olds with 0, 5, and 10 years of experience, respectively, that Plan 1, the pattern plan, seems slightly safer. Even these comparisons are probably within the sampling error.

On an overall, ex ante basis, table 12.12 indicates that Plan 3 is always the safest. These conditional calculations strongly confirm those reported in table 12.7.

Table 12.12 Ratio of Standard Errors to Means: Second Set of Simulations Only Retirement Benefits Are Paid (All Workers).

Experience	Age	Plan 0	Plan 1	Plan 2	Plan 3
0	25	1.2242	1.2068	1.2472	1.1888
	40	.7015	.6926	.7025	.6644
5	30	.9067	.8892	.9133	.8589
	40	.5607	.5494	.5376	.4919
10	35	.7194	.6768	.7011	.6241
	50	.4959	.4549	.4618	.4256
15	40	.6136	.5642	.5864	.5336
	55	.4412	.3862	.3840	.3796
20	45	.5545	.4993	.5194	.4823
	60	.3154	.3453	.2889	.3429

Table 12.13 Ratio of Standard Errors to Means: Second Set of Simulations Only Retirement Benefits Are Paid: Only Retired Workers with Vested Benefits Included

Experience	Age	Plan 0	Plan 1	Plan 2	Plan 3
0	25	.0828	.0657	.1006	.0594
	40	.0876	.0744	.1143	.0808
5	30	.0979	.0800	.1086	.0704
	45	.0986	.0872	.1114	.0908
10	35	.1044	.0879	.1056	.0721
	50	.1280	.1161	.1240	.1189
15	40	.1253	.1082	.1140	.0858
	55	.1698	.1587	.1381	.1584
20	45	.1410	.1257	.1146	.0982
	60	.1983	.2387	.1791	.2345

Comment Alan J. Auerbach

This paper is in some ways easy to discuss but in other ways difficult. Jerry Caren has carefully laid out for us what he has done, so that we may pause at each stage for an evaluation. Ultimately, however, I find it hard to know what message about the riskiness of private pensions I should take away from the simulation results that are presented in the paper. For me, much of the difficulty lies in the concept of risk that Green uses, the different sources of risk in the pensions analyzed, and the many assumptions necessary to generate concrete numerical examples. All in all, the pa-

Alan J. Auerbach is associate professor of economics at the University of Pennsylvania and a research associate of the National Bureau of Economic Research.

per has helped clarify a number of issues that one must confront in assessing the risk inherent in private pension plans, and in the process has shown how difficult a task Green has set for himself. He should be commended for his efforts in exploring this extremely important but little researched question.

The numerical simulations in the paper consider the riskiness of four prototypical private pension plans, as measured by the coefficient of variation of the present value of pension benefits or total compensation, from the viewpoint of a representative worker at one of eight stages in his age and experience with a firm. Simulation is used because of the enormous complexity of the underlying random distribution of pension benefits.

Each worker, identical *ex ante*, faces four sources of uncertainty: (1) date of death; (2) the prospect of preretirement disability; (3) the prospect of preretirement withdrawal from the firm; and (4) the stochastic evolution of wages over time. All sources of risk are assumed to be exogenous from the individual's viewpoint, as well as independent. This rules out, for example, the withdrawal probability's being related to the individual's wage rate or the date of death's being hastened by early disability. These simplifications may be restrictive, but there is little alternative available, given the paucity of data and the great complexity already characterizing the problem.

The age-specific probabilities of death, disability, and withdrawal, presented in table 12.2, are based on actuarial data. The wage process is assumed to be one with an overall stochastic component, as well as individual-specific, or "idiosyncratic," risk. An individual's wage at date t , W_t , may be expressed in terms of his last year's wage by

$$W_t = W_{t-1} (1 + \bar{w}_t + \eta_t),$$

where \bar{w}_t is the overall, or "market," disturbance, itself generated by a Markov process with an autocorrelation coefficient of .9, and η_t is an independent and identically distributed individual error. Both η_t and the white noise component underlying the process for \bar{w}_t are normally distributed, the latter with a nonzero mean intended to capture secular wage growth at an expected rate of 1% per year. Without the term \bar{w}_t , this wage process would be a random walk in terms of the logarithm of the wage. With the highly autocorrelated market component, wherein the growth rate of overall wages is very close to a random walk, actual wage variation over time comes to consist more and more of the market component. Whether this outcome is appropriate is a question Green does not address, nor are we given any reason to have strong faith in the variances chosen for the two sources of risk that generate the wage process. Even if these assumptions come from the best "point estimate" available, a sensitivity analysis would be most enlightening.

The four pension plans considered vary with respect to the method of retirement benefit calculation and the extent to which they provide disability and survivor preretirement death benefits. They also differ with respect to the treatment of early retirement, but early retirement is ignored in the simulations. Vesting provisions are similar with full or partial vesting at 10 years of service with any remaining vesting occurring shortly afterward.

Retirement benefits under Plan 0 (conventional), 2 (career average), and 3 (final salary) are all based on some average over time of an individual's own wages, with the names of the plans reflecting fairly accurately the extent of averaging. Since there is no inflation in the model, one would expect that Plans 0 and 2, which are based on salaries from a certain number of the best years, would be more likely to include wages from early years; that is the case in reality. Plan 1, the pattern plan, provides for a fixed dollar amount multiplied by years of service, but Green interprets this, appropriately I think, as being implicitly indexed to overall wages in the year of retirement. One should note, of course, that this assumption ignores additional risks that would be associated with the presence of noise in the implicit indexing process.

The four plans also differ in their provision of death and disability benefits. Plan 0 provides neither, Plan 1 provides both, Plan 2 provides only for disability, Plan 3 only for death. The spanning of the set of possibilities is either quite fortunate or quite unfortunate, depending on where one's interest lies. I must vote for the latter, if what we seek is to understand the effects of various pension provisions on individual welfare rather than the impact of particular plans themselves. This is a problem that could be remedied by considering hypothetical plans that differed only with respect to benefit calculations or only with respect to auxiliary provisions. At present, I suspect that the differences among the plans in the latter are what generate many of the differences in the results for the plans. This may be inferred from table 12.6, which presents the sample correlations between the present values of wages and benefits for individuals starting at each of the eight points analyzed in the paper. The lowest value, .12, occurs for individuals who are age 60, despite the fact that, for those who reach retirement at 65, pension benefits will be based on the four highest years of own wages, quite possibly four of the remaining five years. One would expect this to lead to a high correlation, but this ignores the fact that states in which death or disability intervene are included in the calculation. While such states have a relatively small chance of occurring, the change in benefits if they do is enormous.

This result foreshadows a problem in separating the effects of particular differences among plans but also brings out another difficulty with the analysis: the symmetric treatment of all states of nature. For example, a

plan with a death benefit equal to half the nominal retirement amount might be just what would be required to insure completely per capita family consumption against a husband's death, yet this plan would be deemed riskier than one with full death benefits. Similar problems are associated with disability. There is a particular problem in the treatment of withdrawal, where all subsequent work experience and pension benefit accruals are ignored. Given the usual backloading of benefit accruals associated with pensions, a person withdrawing to take a better job would, by the current analysis, receive a low benefit relative to wages, even if vested. Finally, since the analysis considers only the risks associated with pensions and wages, we have no sense of which sources of risk are diversifiable and which are not. Even for those workers who are liquidity constrained, social security constitutes an important form of wealth that may offset some of the risks included in private pensions. Indeed, some of the pension provisions may have been designed to take the preexistence of social security into account.

With all of these difficulties in mind, I turn at last to the paper's simulation results. These results were generated for each of the eight cases by following 900 individuals through their working lives until some terminal event occurred. For example, no wage calculations were necessary for those already disabled. The same sample of 900 is used for each of the four pension plans in a given case, so that eight samples were generated overall. While Green argues that this sample size appeared large enough to eliminate most sampling error, this statement does not apply to subsamples as small as 11, conditioned on certain events such as disability.

Because I am unsure how to interpret the measured risk associated with withdrawal, disability, or death, I would find most meaningful calculations based on the condition of a worker actually reaching retirement. The closest Green comes to presenting such results is in table 12.11, which also includes those that withdraw. To focus on retirement, we can look at calculations for those of later ages for whom withdrawal is unlikely, and here there is a counterintuitive result: as in the full-sample calculations in table 12.7, the final salary plan is generally the safest. It is safer than the pattern plan in all but two of the eight cases in the table, even though the pattern plan is independent of the risks of individual wage variation. It is safer than the career average plan in all eight cases, and safer than the conventional plan in seven of eight, despite the fact that the latter two plans should permit more lifetime averaging of wage risks. I find the robustness of these results disturbing, because I can think of no convincing explanation for them. This demonstrates, perhaps, one of the weaknesses of simulation analysis of complicated problems: we cannot simply look at the formula for a derivative to see the origin of an outcome. However, sensitivity analysis is the best available alternative and would be a most helpful addition to the current study in light of the results presented.