

Nicholas Barr and Peter Diamond

The economics of pensions

**Article (Accepted version)
(Refereed)**

Original citation:

Barr, Nicholas and Diamond, Peter (2006) The economics of pensions. [Oxford Review of Economic Policy](#), 22 (1). pp. 15-39

DOI: [10.1093/oxrep/gri002](https://doi.org/10.1093/oxrep/gri002)

© 2006 [Oxford University Press](#)

This version available at: <http://eprints.lse.ac.uk/2630/>

Available in LSE Research Online: January 2008

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final manuscript accepted version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this version and the published version may remain. You are advised to consult the publisher's version if you wish to cite from it.

The economics of pensions

Nicholas Barr and Peter Diamond

January 2006

Oxford Review of Economics Policy, Vol. 22. No. 1, Pensions, forthcoming

Draft. Not for quotation

The economics of pensions¹

Nicholas Barr² and Peter Diamond³

Abstract

This paper sets out the economic analytics of pensions. After introductory discussion, successive sections consider the effects of different pension arrangements on labour markets, on national savings and growth, and on the distribution of burdens and benefits. These areas are controversial and politically highly salient. While open about expressing our own views, the main purpose of the paper is to set out the analytical process by which we reach them, to enable readers to form their own conclusions.

1 The backdrop

This paper has a two-fold purpose. It sets out the economic analytics of pensions without discussion of empirical magnitudes and outside the context of any particular country, with the intention of giving readers a systematic way of thinking about the topic. The paper is also intended as a contribution to a continuing debate, hence part of the discussion rebuts arguments that we regard as false, or equivocal, or true in some circumstances but not necessarily always. Specifically, we argue that much analysis is incomplete and over-simplified: focussing on one objective while ignoring others; assuming an idealised economy with well-informed agents and no distortions such as taxes and missing markets; comparing one steady state with another, when the underlying issue is a move from one steady state to a different one; or ignoring distributional effects.

The opening section sets out some background matters: the objectives of pension systems, types of pension arrangement, and the economics of pensions. Sections 2, 3 and 4 discuss in turn pensions and labour markets (mainly microeconomic), finance and funding (mainly macroeconomic), and distributional issues.

1.1 The objectives of pension systems

From an individual viewpoint, income security in old age requires two types of instruments: a mechanism for consumption smoothing, and a means of insurance.

¹ This paper is a shortened version of Barr and Diamond (forthcoming, Ch. 2), which is an outgrowth of our participation in a panel examining pension issues in China. We are grateful to other panel members and commentators on that report and for helpful comments on this paper from Christopher Bliss, Robert Hancké, Dieter Helm, David Hendry, Stephen Nickell and members of the OXREP editorial board at a seminar in Oxford.

² Professor of Public Economics, London School of Economics and Political Science; Email: N.Barr@lse.ac.uk.

³ Institute Professor, Massachusetts Institute of Technology; Email: pdiamond@mit.edu.

Consumption smoothing

People seek to maximise their well-being not at a single point in time, but over time. Someone who saves does so not because extra consumption today has no value, but because he values extra consumption in the future more highly than extra consumption today. A teenager who saves for a flight ticket is making a judgment that she will get more enjoyment from the trip than from spending the money now. Similarly, most people hope to live long enough to be able to retire. Thus a central purpose of retirement pensions is consumption smoothing – a process which enables a person to transfer consumption from her productive middle years to her retired years, allowing her to choose her preferred time path of consumption over working and retired life.⁴

Insurance

In a model of certainty, individuals save during their working life to finance their retirement. Important in the case of pensions is that people face a range of uncertainties, including how long they are going to live. Thus a pension based on individual saving faces the person with the risk of outliving those savings, or of consuming very little to prevent that happening. Though any one person does not know how long he is going to live, the life expectancy of a large group of people is better known. Thus in principle, the members of the group could agree to pool their pension savings, with each person drawing a pension based on (a) the group's life expectancy and (b) the total amount he or she had contributed to the pool. In addition, members of the group could pay others to absorb the longevity risk.

This is the essence of annuities, whereby an individual exchanges his pension accumulation at retirement for regular payments for the rest of his or her life, thus allowing people to insure against the risk of outliving their pension savings. Pension systems can also protect spouses and young children should a worker die before retirement, and can insure against disability.

Are voluntary arrangements sufficient? In the simplest of all worlds a person provides for his pension through voluntary savings to achieve his optimal time path of consumption and through an annuity to protect himself (and his spouse) against the longevity risk. Were matters that simple, pensions could be left to voluntary decisions and private insurance, with no need for government involvement. There are two sets of reasons why this approach, on its own, is insufficient. First, the simple model assumes that there is perfect information (apart from date of death) and no other distortions. These assumptions are useful to formulate a simple theory, but bad guides to policy design in a world with imperfect information, missing markets, risk and uncertainty, and distortions such as progressive

⁴ This process is explained more formally by the simple Fisher model. See, for example, Barr (2001, Ch. 2).

taxation. Moreover, there are serious concerns about the abilities of individuals to make the most of the market opportunities available to them. The simple models, in implicitly assuming a first-best world, ignore a range of market failures, and thus assume away the very problems that government intervention is designed to address. In contrast, second-best analysis seeks the optimal policy given the presence of such distortions.

A second reason for government involvement is that public policy generally has objectives additional to improving consumption smoothing and insurance, notably poverty relief and redistribution.

Poverty relief

Poverty relief targets resources on people who are poor on a lifetime basis, and thus unable to save enough. As a practical matter, poverty relief also has to address transient poverty. Such programmes can target all the elderly or may concentrate on those who have contributed to the pension system.

Redistribution

Pension systems can redistribute incomes on a lifetime basis, complementing the role of progressive taxes on annual income. Lifetime redistribution can be achieved by paying pensions to low earners that are a higher percentage of their previous earnings (i.e. a higher replacement rate), thus subsidising the consumption smoothing of lower earners. Since life-long earnings are uncertain from the perspective of an individual, such a system provides some insurance against low earnings. There can also be redistribution towards families, for example paying a higher pension to a married couple than to a single person, even though both families have paid the same contributions.

Pension systems can also redistribute across generations, for example if a government reduces the contribution rate of the present generation, thereby requiring future generations to pay higher contributions or have lower pensions.

Other objectives

Alongside these primary objectives, pensions policy may have secondary goals, including economic development broadly and economic growth specifically. Badly-designed pensions may create adverse labour-market incentives. Excessive public pension spending contributes to high tax rates, putting growth at risk. Conversely, pension arrangements can assist the operation of labour and capital markets and may encourage saving. There is debate about the relative weights accorded to old age security and to these secondary objectives.

1.2 Types of pension

Pensions can be arranged in different ways, relating to (a) the way they are organised and (b) the relation between contributions and benefits.

Fully-funded, partially-funded and pay-as-you-go pensions

In a fully-funded scheme, pensions are paid out of a fund built over a period of years from its members' contributions. With Pay-As-You-Go (PAYG) schemes, in contrast, pensions are paid out of current income. While we describe the polar cases, partial funding represents a continuum between them.

FULLY-FUNDED SCHEMES are based on savings – contributions are invested in financial (or possibly physical) assets, the return on which is credited to the scheme's fund. Funding is thus a method of accumulating financial assets, which are exchanged for goods at some later date. While fully-funded schemes can take many forms, in principle they always have sufficient reserves to pay all outstanding financial liabilities (or equivalently, liabilities are defined by available funds).

If there is no redistribution across generations, a generation is constrained by its own past savings and a representative individual gets out of a funded scheme no more than he has put in.⁵ If, in addition, there is no direct redistribution across individuals, when an individual retires, the pension fund will be holding his past contributions, together with the interest and dividends earned on them. This accumulation finances the person's consumption in retirement, through an annuity or in some other way.

PAYG SCHEMES are usually run by the state. They are contractarian in nature, based on the fact that the state can but does not have to accumulate assets in anticipation of future pension claims, but can tax the working population to pay the pensions of the retired generation. Most state pension schemes are primarily PAYG.⁶

From an economic viewpoint, PAYG can be looked at in several ways. As an individual contributor, a worker's claim to a pension is based on a promise from the state that, if he pays contributions now, he will be given a pension in the future. The terms of the promise are fairly precise, being set out in each country's social security legislation (although

⁵ In reality, matters are more complex: real rates of return, and thus future assets, are a random variable (see, for example, Burtless, 2002); analogously, future liabilities are a random variable, particularly if life expectancy is uncertain. Thus analysis in terms of simple present values, though useful conceptually, is not always a good guide to policy.

⁶ PAYG schemes have also been run by corporations. Just as a state PAYG scheme is dependent on the presence of a future tax base, so a corporate PAYG scheme is dependent on the presence of future corporate earnings to pay pensions. Because of the risk of non-payment, such schemes have been found unsatisfactory and banned in many countries.

subject to legislative change).⁷ From an aggregate viewpoint, the state is simply taxing one group of individuals and transferring the revenues to another, whether viewed on an annual or a lifetime basis. State-run PAYG schemes, from this macroeconomic perspective, are little different from other income transfers, although the determinants of who pays and who receives and the incentive structure can be very different from other income transfer systems.

A major implication of a PAYG system is that it relaxes the constraint that the benefits received by any generation must be matched by its own contributions. Samuelson (1958) showed that with a PAYG scheme it is possible in principle for every generation to receive more in pensions than it paid in contributions, provided that the rate of growth of total real earnings exceeds the interest rate indefinitely; this can happen when there is technological progress and/or steady population growth and excessive capital accumulation (see Aaron 1966). Since this does not appear to be empirically relevant over the longer term, the real role of PAYG is to redistribute across generations and to share risks across generations.

DEBATES. There is considerable controversy over the relative merits of PAYG and funded schemes.⁸ There are debates:

- about the right basic economic model, for example how to model individual behaviour;
- about empirical magnitudes, for example, about labour supply elasticities, and about life expectancy in 2050;
- about the extent of a country's institutional capacity;
- about the political economy of reform, for example, whether citizens regard their pension as safer based on a promise by government or as the owners of capital;
- about ideology, for example about the role of the state, or about the relative weights given to different objectives, in particular the relative weights accorded to poverty relief and consumption smoothing.

The relation between contributions and benefits

Whether funded or PAYG, a separate question is how closely pension benefits are related to a worker's previous contributions. Three approaches are common.

⁷ The nature of the 'promise' can be complex: with incomplete specification of all the circumstances that may occur in the future, the outcome of the promise is dependent on future actions by the promisor. The promisee may or may not be aware of the dependence of outcomes on such future actions or the nature of the future process that the promisee will follow. Fully-funded schemes are also subject to legislative change in the taxation of assets, returns on assets and payment of benefits.

⁸ See Barr (2000), Diamond (2004), Diamond and Orszag (2005) and, for contrasting views, Feldstein (2005) and Holzmann and Hinz (2005). For an attempt to summarise the core of the dispute, see Barr and Rutkowski (2005).

DEFINED-CONTRIBUTION SCHEMES. In a defined-contribution (DC) scheme, also called funded individual accounts, each member pays into an account a fixed fraction of his or her earnings. These contributions are used to purchase assets, which are accumulated in the account as are the returns earned by those assets. When the pension starts, the assets in the account finance post-retirement consumption through an annuity or in some other way. In a pure DC scheme (i.e. one with no redistribution across individual accumulations), a person's consumption in retirement, given life expectancy and the rate of interest, is determined by the size of his or her lifetime pension accumulation, preserving the individual character of a person's lifetime budget constraint.⁹

Though annuities protect the individual against the risks associated with longevity, a pure DC scheme leaves him or her facing the wide range of risks, discussed below, associated with varying real rates of return to pension assets, the risks of future earnings trajectories and the future pricing of annuities. The pure case can be modified to share risks somewhat more broadly, for example via a guaranteed minimum pension or pooling a part of contributions or a legislated response to capital market outcomes. Labour market incentives are affected by the details of asset accumulation, redistribution across accounts and the benefit formula.

DEFINED-BENEFIT SCHEMES. In a defined-benefit (DB) scheme, a worker's pension is based not on his accumulation, but on his wage history, possibly including length of service. A key design feature is the way wages enter the benefit formula. In a final-salary scheme, pensions are based on a person's wage in his or her final year, or few years. Alternatively, the pension can be based on a person's real or relative wages over an extended period, including an entire career. In either case, a person's annuity can be, in effect, wage indexed until retirement. The worker's contribution is generally a fraction of his or her wage; thus the sponsor's contribution is conceptually the endogenous variable in ensuring the scheme's financial balance. DB schemes can have assets held in a central pool.

DB schemes can be run by the state or by employers. Where a state scheme is financed from contributions, the risk of adverse outcomes falls on current contributors; where there is a taxpayer subsidy, the risk falls on taxpayers. In practice, governments change benefits as well as contributions when revenue and expenditure do not balance. Such adjustment can be automatic (indexed) or the result of specifically legislated change.

In an employer scheme, the risk of varying rates of return to pension assets falls on the employer, and hence on some combination of the industry's current workers (through effects on wage rates), its shareholders and the taxpayer (through effects on profits), its

⁹ A structure with funded individual accounts can have redistribution across workers' accounts or from general revenues to the accounts.

customers (through effects on prices) and/or its past or future workers, if the company uses surpluses from some periods to boost pensions in others or modifies the benefit formula relative to expectations. In a pure DB scheme, therefore, none of the risks fall directly on pensioners. In practice, however, company DB schemes may also adjust current and/or future benefits in the light of financial outcomes. A key difference between DB and DC pensions is how and how widely risks are shared.

NOTIONAL DEFINED-CONTRIBUTION (NDC) SCHEMES. A recent innovation internationally, pure NDC systems are conceptually similar to pure DC pensions in the way one aspect of risk is shared, with all adjustment taking place on the benefits side, but different, in that they are not fully funded and may be entirely PAYG. NDC schemes parallel DC pensions:

- Each worker pays a contribution of $x\%$ of his or her earnings, which is credited to a notional individual account, that is, the state ‘pretends’ that there is an accumulation of financial assets.
- The cumulative contents of the account are credited with a notional interest rate, specified by the government, and chosen to reflect what can be afforded.
- At retirement, the value of the person’s notional accumulation is converted into an annuity in a way that mimics actuarial principles, inasmuch as the present value of a person’s benefits (given mortality rates based on the worker’s birth cohort and age) is equal to the value of the person’s notional accumulation, using the notional interest rate as the discount rate.
- The account balance is for record keeping only, because the scheme does not own matching funds invested in the financial market. This explains the term ‘notional’.¹⁰

Thus NDC pensions mimic funded defined-contribution schemes by paying an income stream whose present value over the person’s expected remaining lifetime equals his/her accumulation at retirement, but with an interest rate set by government rules, not market returns. As with DC pensions, there are multiple ways of incorporating a redistributive element in the accounts, including a minimum pension guarantee or by subsidising the contributions of people who are out of the labour force because they are bringing up young children or are unemployed. For fuller discussion see Palmer (2005).

On the face of it, NDC schemes, where benefits depend on a history of contributions are very different from standard DB schemes, where benefits depend on a history of earnings. If contribution rates do not change, however, this distinction is irrelevant, and an NDC scheme can be viewed as a DB scheme with a particular structure of automatic adjustments for demographic and economic realizations. Indeed, an NDC scheme is quite close to some

¹⁰ Preserving the acronym, these have also been referred to as nonfinancial defined contribution schemes.

schemes described in a DB vocabulary, so that the difference between the two approaches should not be exaggerated. More generally, the choice of vocabulary can have political implications for the process of pension reform.

1.3 The economics of pensions

Simple economics

It assists analysis to have three propositions in mind: what matters is output; imperfect information and imperfect decision making are pervasive; and pension schemes face large and unpredictable risks. A fourth important point is that pension arrangements have administrative costs that can be significant.

WHAT MATTERS IS OUTPUT. There are two (and only two) ways of seeking security in old age. One is to store current production for future use. But, housing excepted, this approach is inadequate for most consumption needs: it is expensive; it does not address uncertainty (e.g. about how a person's tastes might change); and it cannot be applied to services deriving from human capital, notably medical services.

The alternative is for individuals to exchange current production when younger for a claim on future production when older. There are two broad ways to do so: by saving part of his wages a worker could build up a pile of *assets* which he would exchange for goods produced by younger people after his retirement; or he could obtain a *promise* – from his children, his employer, or government – that he would be given goods produced by younger workers after his retirement. The two main ways of organising pensions broadly parallel these two types of claim. Funded schemes are based on accumulations of financial assets, Pay-As-You-Go (PAYG) schemes on promises.

The purpose of pensions is to allow people to continue to consume after they have stopped working. Pensioners are not interested in money, but in consumption—food, clothing, heating, medical services. Consumption comes from goods produced at the time – and therefore by younger workers. To that end, future output is central. PAYG and funding are simply financial mechanisms for organising claims on that future output. In macroeconomic terms, though there are differences between the two approaches, those differences should not be exaggerated.

The centrality of output remains true in an open economy. In principle, pensioners are not constrained to consumption of domestically-produced goods, but can consume goods made abroad so long as they can organise a claim on those goods. If British workers use some of their savings to buy Australian factories, they can in retirement sell their share of the factory's output for Australian money to buy Australian goods, which they then import to the

UK. Though useful, this approach is no panacea. The policy breaks down if Australian workers all retire; thus the age structure of the population in the destination of foreign investment matters. Second, if large numbers of British pensioners exchange Australian dollars for other currencies, the Australian exchange rate might fall, reducing the real value of the pension. Thus the ideal country in which to invest has a young population *and* products one wants to buy *and* political and financial stability *and* is large enough to absorb the savings of other countries with ageing populations. Countries with ageing populations include all of the OECD, and many others, China being a notable example.

IMPERFECT CONSUMER INFORMATION AND DECISION-MAKING ARE PERVASIVE. On the microeconomic side, the advantages of consumer sovereignty are predicated on well-informed consumers, a very strong assumption in the case of pensions.

Individuals are imperfectly informed, first, because of uncertainty about the future – individuals are not well-informed because nobody is well-informed. Second, they are imperfectly informed in the face of risk, discussed below.

A third type of imperfect information can arise with complex products like defined-contribution pensions, which are based on an array of financial institutions and financial instruments. Even in the USA there is considerable ignorance. Orszag and Stiglitz (2001, p. 37) quote the Chairman of the Securities and Exchange Commission as stating that over 50 per cent of Americans did not know the difference between a stock and a bond. The problem has equity as well as efficiency implications, since the people who are worst-informed are disproportionately the least well-off. Defined benefit schemes can also be complex and incompletely understood.

For some purposes it is useful to distinguish a fourth type of problem – what New (1999) calls an information-processing problem. An information problem can be resolved by providing the necessary information, e.g. the characteristics of different automobiles, after which the individual can make his own choices. With an information-processing problem, in contrast, matters are too complex for agents to make rational choices even if the necessary information is provided. The problem can arise (a) where the time horizon is long, as with pensions, (b) where the good or service involves complex probabilities, including, for example, life expectancy (the failure in this case is an inability to process probabilities), (c) where the information is inherently complex, as with complicated pension products or (d) where the use of the information requires sophisticated analysis.

Some ignorance – information problems – can be reduced by public education. However, some is inherent and cannot be resolved in that way. Even financial sophisticates cannot necessarily be regarded as well-informed consumers. Given the high potential cost of

mistaken choice, imperfect information creates an efficiency justification for stringent regulation to protect consumers in an area where they are not well-enough informed to protect themselves. Evidence suggests it is difficult and expensive to provide information that succeeds in altering behaviour.

Beyond imperfect information acquisition and processing there are issues of the quality of decision making as it affects both workers and their families. If workers are not time-consistent in their decisions about savings or annuitization and if their decisions do not pay sufficient attention to the future needs of other family members, then there are bases for state interventions, bases that have been recognized for centuries, for example through restrictions on estates to protect widows.

PENSION SCHEMES FACE LARGE RISKS THAT ARE HARD TO PREDICT.

- Macroeconomic shocks affect output, prices or both.
- Demographic shocks affect all pension schemes (see section 3.2), by affecting market prices and quantities and pension claims.
- Political risks affect all pension schemes because all depend critically – albeit in different ways – on effective government.
- Management risk can arise through incompetence or fraud, which imperfectly-informed consumers generally cannot monitor effectively.
- Investment risk: private and public pension accumulations held in the stock market until retirement are vulnerable to market fluctuations.
- Annuities market risk: for a given pension accumulation, the value of an annuity depends on remaining life expectancy and on the rate of return the insurance company can expect over those years (and is thus also a form of investment risk).

Private insurance markets can help individuals to bear some of the risks inherent in preparing for retirement. But there are limits to private insurance from adverse selection, from selling costs, from the limited ability of consumers to make good decisions, and from incomplete markets for risk-sharing, particularly across cohorts. With social insurance, as discussed in section 4, the intention is for risk to be shared more broadly. The costs of adverse outcomes can be borne by the pensioner through lower pensions, by workers through higher contributions, by the taxpayer, through tax-funded subsidies to pensions, and/or by future taxpayers and beneficiaries if subsidies are financed by government borrowing.

COSTS. The previous arguments all apply even in a frictionless world. But analysis must also take into account the fact that any method of arranging for future consumption has administrative costs. These include the costs of record-keeping and the costs of transactions

insofar as there are accumulations of assets or purchases of benefit streams. Different ways of organizing future consumption have very different costs and thus provide very different levels of future consumption. For example, the individual mutual fund market is far more expensive than the institutional mutual fund market.

But not too simple

This paper offers analysis and conclusions relating to three sets of issues.

PENSIONS AND THE LABOUR MARKET. It is not possible to have a modern economy without distorting the labour market. Analysis of pension systems has to recognise the tradeoff between efficiency in labour markets, on the one hand, and the contribution to the goals of consumption smoothing, insurance, redistribution and poverty relief, on the other.¹¹ Thus the real issue is to balance distortions with other goals, not to pretend that there is a way to accomplish multiple goals without distortions. What is needed, therefore, is second-best analysis, which considers the impact of the entire programme for retirement income. It can lead to error to consider one part in isolation. These topics are discussed in section 2.

PENSIONS AND NATIONAL SAVINGS. A mandatory retirement income system affects national savings. An important issue, therefore, is the extent to which the system is funded in a way that increases national savings, and so increases future output. Section 3 presents a framework for thinking about the extent of funding – recognising that, depending on savings needs, a good system can have any degree of funding, from none to full.

DISTRIBUTIONAL EFFECTS are discussed in section 4. While private insurance markets, along with capital markets, are devices for sharing risks, a public pension system can improve risk-sharing in ways that are not available to the market. Furthermore private insurance markets are subject to some significant limitations that public provision can overcome. Section 4 also discusses how the design of the pension system affects lifetime income distributions.

Common errors

A number of common errors are discussed as they arise.

CONSIDERING ONE OBJECTIVE IN ISOLATION. Pensions have multiple objectives; some writers focus on one, ignoring the others. It is right to debate the relative weights accorded the various objectives, but policy analysis that focuses on a single objective, particularly if it does so implicitly, will be flawed.

¹¹ While the vocabulary of a tradeoff between labour market efficiency and other goals is useful, it is worth remembering that with incomplete markets and imperfect decision-making there are interventions that can improve labour market efficiency (suitably defined) while also advancing other goals.

IMPROPER USE OF FIRST-BEST ANALYSIS. It is mistaken to focus on labour market distortions while ignoring or downplaying the contributions to the various goals of pension systems – contributions that are not available without distortions. An error discussed below is to argue that actuarial benefits are optimal because they minimise distortions. Of course, one should not ignore distortions, nor design pensions in ways that have larger distortions than are justified by sufficient contribution to goals, but it does not follow that minimising distortions is necessarily the right objective.

IMPROPER USE OF STEADY STATE ANALYSIS. It is mistaken to focus on a pension system in steady state, while ignoring or underplaying the necessary transition steps to get from one steady state to another. This is particularly an issue when considering funding of pensions. An error discussed below is to argue that funding is necessarily superior because stock market returns are higher than the rate of growth of the wage base, which determines the return to PAYG schemes.

IGNORING DISTRIBUTIONAL EFFECTS. Pension systems can redistribute across cohorts of different birth years, so one needs to consider both those who gain from enlarging pensions and those who lose because of the need to finance the pensions at some time, possibly in the future. It is mistaken to look at one while ignoring the other.

2 Pensions and labour supply

This section discusses the influence of pension design on labour markets during working life and when considering retirement.

2.1 The effects of benefit design during working life

Two issues stand out: the problems associated with final salary schemes, and the extent to which an actuarial relationship between contributions and benefits is or is not an advantage.

Problems with final salary schemes

Corporations use pensions to attract and retain workers. Historically, many schemes paid pensions at a standard retirement age that depended on length of service and the worker's wage towards the end of his or her career. Such a structure makes it easy for workers to see the advantages of staying with the firm until retirement.

Schemes of this sort can create labour-market problems. A young worker recognises that current earnings do not affect the size of his future pension, weakening the incentive to work extra hours or to take on a harder job at higher pay. There are, of course, offsetting incentives, since the worker might recognise that hard work and accomplishment improve the chances for promotion, and so a higher wage and hence a higher pension later on.

The opposite incentive operates towards the end of a career, where workers might be over-eager to work extra hours. An extreme version of this problem arose in Boston, where the subway system bases pensions on the earnings (not the base pay) of workers at the end of their careers. Older workers therefore do a great deal of the overtime, which has caused accidents when, working excessive hours, they fall asleep at the controls of trains. One need not go so far as endangering lives to see that such pension arrangements can have adverse incentives. A similar problem in a large organisation is promotion toward the end of a career to raise the pension entitlement of a person favoured by the middle managers who control promotions. The shorter the period of earnings used in determining benefits, the stronger the incentive for such manipulative collaboration.

A third problem with corporate final-salary DB schemes is their effect in locking a worker into employment with that corporation.¹² Historically, indeed, that was one of the main purposes of that benefit design. In a modern economy, the efficiency costs of the resulting impediments to labour mobility are likely to be substantial.

A fourth problem relates to the distribution of pension incomes, which favours workers whose earnings rise more rapidly, particularly towards the end of a career. Since highly-paid workers tend to have more rapidly rising earnings, the system favours the best-off. This can be regarded as unfair.

For all these reasons, in corporate or industry schemes a worker's pension should depend on most or all of his or her earnings history.

National schemes face the same problems for the same reasons. Within a single corporation, these effects can be lessened by the other controls a corporation has in relation to its workers and by relating pensions to base pay not actual earnings. In a national system, the government does not have similar controls over the entire economy. Thus, it is important that a national system bases benefits on most or all of a worker's earnings history once the social security administration has the necessary administrative capacity.

Problems with strict adherence to actuarial benefits

It is frequently argued that a strictly actuarial relationship between contributions and benefits is optimal: 'Funded defined-contribution schemes are the closest to an actuarially fair system, so the labour market distortions should be low' (Holzmann and Hinz, 2005, p. 50). There are three sets of arguments: that actuarial benefits minimise distortions to labour supply, that they improve compliance with contribution conditions, and that they encourage later retirement.

¹² For a simple example, see Barr (2001, p. 153).

In a first-best world, actuarial benefits face each individual with an efficient choice between consumption when younger and consumption when older. In practice, however, policy design must address at least three sets of market imperfections. People can be myopic and/or imperfectly informed, giving a justification for compulsion. The problem is non-trivial, and means that the simple assumption of rational utility maximisation may not hold, particularly in the face of the major information problems discussed earlier.¹³ A second problem is missing markets. The market for indexed contracts, for example, is thin. Insurance with asymmetric information requires distorting labour market decisions in order to have insurance.¹⁴ Progressive taxation is a third deviation from first-best. Diamond argues that in the comparison between defined contribution and defined benefit schemes, ‘there is no simple dominance of one over the other in the presence of other labour market distortions’ (2002, p. 57). Formulating the issue as an optimal taxation problem would make it clear that in a second-best world, a scheme that is strictly actuarial is not, in general, efficient.

What, then, of the three specific advantages claimed for actuarial benefits? Actuarial benefits will generally not minimise labour market distortions, given the presence of other distortions. Furthermore, pensions have objectives additional to consumption smoothing, for example poverty relief, and, as already noted, the policies (e.g. taxation) necessary to achieve those other objectives inescapably involve labour market distortions. And third, the provision of insurance against adverse labour market outcomes, particularly towards the end of a career, calls for deviations from actuarial insurance to provide better protection, given the asymmetry of information on the extent to which low labour market participation is due to choice (preference for more leisure) or constraints (low pay or no work available). In sum, (a) actuarial benefits do not minimise labour-market distortions and (b) minimising labour-market distortions is not the right objective – policy has to balance labour-market efficiency against the various objectives of pension schemes.

Actuarial benefits may improve compliance with contribution conditions where individuals are well-informed and not liquidity constrained (that is, can borrow at a market or near-market interest rate). In reality, people may be badly-informed about the relation between contributions today and pensions tomorrow; they may be myopic; or constraints on their borrowing capacity might lead them to choose current over future consumption. For all these reasons, low compliance in a country like Chile, with largely actuarial benefits, is a major concern, as discussed in the paper (this journal) by Arenas and Mesa Lago.

¹³ Mandatory saving for retirement also affects labour supply. But just as workers may not perceive the advantages of saving adequately for retirement, they may not properly link today’s work with future benefits.

¹⁴ Indeed the theorem is that an absence of labor market distortions in the presence of asymmetric information is a sign of nonoptimal provision of insurance.

Actuarial benefits, it is argued, encourage the appropriate amount of later retirement by offering a person who defers retirement a pension that is sufficiently larger as to involve no implicit tax.¹⁵ As above, for reasons of insurance or to offset poor decisions, a zero implicit tax may not be optimal. Moreover, a DB system could adjust benefits for work beyond the earliest pensionable age on an actuarial basis even though the determination of initial benefits at the earliest pensionable age is not actuarial. This is discussed further below.

In all three cases – minimising distortions, improving compliance, and encouraging later retirement – the simple argument holds only in a first-best world. That does not mean that the relation between contributions and benefits is unimportant. Indeed, good policy design should avoid obvious and major distortions in the relationship between contributions and benefits. But it does mean that a strictly actuarial relationship is generally sub-optimal.

2.2 Determining benefits at retirement

An array of design features at the time a person retires can have major effects on labour markets. This section considers in turn the relation between retirement age and aggregate unemployment; the undesirability of mandatory retirement; issues surrounding the choice of age at which a worker is first entitled to benefit; adjusting benefits where a person retires later; and adjusting pension schemes as life expectancy increases.

Retirement age and unemployment

The common view that early retirement eases unemployment is generally mistaken. From a long historic perspective, developed countries have seen a vast decrease in the average retirement age, yet unemployment has shown no trend decrease. Evidence for a number of countries over a 10-year period shows no pattern whereby countries which encourage early retirement have lower unemployment. Indeed, when comparing two countries, it is possible to observe one with higher unemployment over an extended period, and then the other, even where their retirement systems do not change significantly.

It is mistaken for several reasons to think in terms of a fixed number of jobs. First, increased numbers of workers, by exerting downward pressure on wages and by making it easier to find suitable workers, tend to encourage the creation of new jobs. Thus the number of jobs is variable, and is influenced by the number of workers available. Second, early retirement frequently does not remove workers from the labour force, since some workers continue to work while receiving a pension. Third, in a developing economy urban unemployment depends on migration as well as on the availability of jobs. Any attempt to reduce urban unemployment by encouraging early retirement may be dwarfed by migration.

¹⁵ This ignores the use of uniform annuity pricing on workers with different life expectancies, as is inevitable.

Thus it is mistaken to allow or mandate early retirement (which is long term) as a palliative response to unemployment which is generally short term. Better to focus on incentives which encourage long-run growth than to distort the labour market in the vain hope that retirement will have a large impact on unemployment. Similarly, disability benefits should be awarded on the basis of disability, not as a response to unemployment.

Mandatory retirement

Forcing people to leave the labour force has no sustained benefit for workers seeking jobs. Thus there is no reason to have a mandatory retirement age on a nationwide basis. Older workers differ greatly in their health, interest in work, ability to work, and job opportunities. Employers differ greatly in their potential for and need for older workers. Flexibility in ending employment relationships is an important part of the efficient long-run use of labour. The USA, with some exceptions, has made mandatory retirement illegal *at the firm level*, and the EU is following suit. But it is not necessary to go as far as this to recognise a role for allowing firms and workers to select retirement ages. Mandatory retirement on a nationwide basis is neither necessary nor desirable.

What 'retirement age'?

As discussed by Banks and Smith (this journal), the concept of retirement is multi-dimensional. When thinking about 'the retirement age' in a pension scheme, two variables are particularly important. Corporate schemes often give a single retirement age a central role, perhaps with a smaller or larger pension for earlier or later retirement. We refer to the age that plays this central role as the age for full benefits. For a national system, it may be more useful to think in terms of (a) the earliest age at which a worker is allowed to start benefits (earliest pensionable age or earliest eligibility age – EEA) and (b) the increase in pension of someone who delays the start of benefits beyond that age. Different countries have different EEAs. The USA has both an earliest pensionable age (62) and an age for full benefit (65, rising to 67). In the UK, the two ages are the same.

What factors should guide the choice of earliest eligibility age? Raising the EEA does not help long-run pension finance if benefits are actuarially adjusted for the increase in the age at which they start. Raising the age at which benefits start lowers costs only if accompanied by a decrease in the level of benefits at each age below what it would have been under the old system.

Increasing the EEA helps some workers and hurts others. The age should be chosen to balance these gains and losses at the margin. Increasing the EEA from 65 to 66:

- Hurts workers who should sensibly stop working at 65 but do not have enough savings to stop working without a pension.
- Helps workers who ought to wait until 66 but who would, given the choice, retire at 65 on a pension that may be inadequate as the worker, and possibly spouse, age.
- Helps workers who retire at 65 and can afford to live from savings until benefits start at 66, by providing higher pensions and so more insurance.

An optimal EEA strikes a balance between helping some workers and hurting others – it should be set in the interior of the range of sensible retirement ages for different workers. Whatever the EEA, the system should be designed to allow flexibility in retirement decisions.

Benefit levels and later retirement

VARIATION IN RETIREMENT AGE. Having selected an age for full benefits, a traditional corporate scheme has a formula that calculates a person's pension at the age for full benefits as a function of (a) years of service and (b) the person's earnings in the years relevant to the benefit formula. However, some firms want some workers to continue beyond the age for full benefits, at least on a part-time basis, and in other cases it may be in the interests of both worker and employer for the worker to retire at a younger age. Actuaries can estimate what reduction for earlier retirement allows a firm broadly to break even from offering an early retirement option.

Actuarially fair adjustments, however, may or may not be in the firm's best interests. A firm might want to give workers more or less encouragement to retire early by setting pensions above or below the level that would break even. Setting pension levels for these alternative options represents an additional control variable for encouraging or discouraging retirement at different ages, separate from the rules that determine benefits at the age for full benefits. In addition, a firm can choose at which age the early retirement option is available, and may offer that opportunity only to a subset of its workers.

If a firm wants to retain some workers beyond the age for full benefits, it can offer a larger pension for delayed retirement. Alternatively, a firm could pay benefits, wholly or in part, while continuing to employ the worker, for example by hiring the worker as a consultant after he or she has formally retired and started to receive benefits. Firms recognise that they do not want all their workers to retire at exactly the same age, recognising differences in jobs and in the different abilities of different workers.

Similar issues arise in a national system. Whatever the rules for pensions at some normal age, there are good reasons – for the economy and for society – for different workers to retire at different ages. Some workers enjoy their work and want to continue working. Others no longer enjoy their work (if they ever did) and want to stop as soon as they can

afford a decent retirement. A good pension system will not excessively discourage the first group from continuing to work at ages at which the second group will already have retired.

BENEFIT ELIGIBILITY AND WORK AT DIFFERENT AGES. Typically, retirement from a firm is a condition for the start of a pension. But that does not necessarily mean the end of all work. Many workers retire from one firm, collect a corporate pension, and then work elsewhere; and some firms allow workers who have retired from full-time work to continue part time, with access to some or all of their pension.¹⁶ As discussed, such flexibility is appropriate.

National systems can also choose the links between continued work and receipt of benefits. Pensions can start at a given age only where a worker has stopped working totally (or has low earnings), or whether the worker stops work or not. Or there can be an age-varying rule – a range of ages where benefits are paid only if work stops, after which benefits are paid irrespective of a person’s labour market activities.¹⁷ In addition, a worker who is eligible for a pension might be allowed to defer benefits so as to have larger benefits once he or she does start.¹⁸

In sum, there are two elements to the relation between pension benefits and age at which pension is first received:

- The pension should be larger for a worker who is older when benefits begin in order to preserve incentives to work until a suitable age for stopping work.
- Benefits should start at a given age without requiring an end to work, or they should increase significantly for a delayed start.

Adjusting pensions for longer life expectancy

How should a pension system be adjusted to reflect differences across cohorts? Specifically, how should contributions and benefits vary, and how should the EEA and the adjustments for early and late retirement vary?

Although mortality rates are likely to continue to decline there is debate about the speed of change. History suggests that we should expect significant deviations in the future from current projections, even if those projections are on average accurate over long periods. In 1981, the UK Government Actuary projected that male life expectancy at 65 in 2004

¹⁶ From April 2006, UK legislation allows a worker to receive an occupational pension while continuing to work for the same employer.

¹⁷ In the USA, benefits are subject to an earnings test between age 62 and the Normal Retirement Age (i.e. the Age for Full Benefits), which is in the process of changing from 65 to 67. There is no earnings test after the Normal Retirement Age.

¹⁸ The UK has no earnings test at State Pensionable Age, but workers receive a higher pension if they choose to delay the initial receipt of benefit.

would be 14.8 years; in reality, it was 19 years, a 28 per cent error. Thus the Second Report of the UK Pensions Commission notes that,

‘...around the 2003-based [Government Actuary’s Department] principal projection of life expectancy for a man aged 65 in 2050 of 21.7 there was a wide and asymmetric range of uncertainty stretching at least from 20.0 to 29.0, but with small probabilities of still wider divergence.... It is therefore essential that both state pension policy and occupational pension provision, in both the public and private sectors, is designed to be robust not just in the face of increasing life expectancy but of major uncertainty about how fast that increase will proceed (2005, p. 90).

Thus projected mortality improvements face a widening funnel of doubt about future outcomes. If current legislation sets future adjustment factors, they will generally not match actual mortality rates. It is, of course, always possible to change the adjustment factors. But legislating change may be difficult and may include a slow transition. Thus there is considerable advantage in designing a system which can, at least up to a point, respond automatically as uncertain outcomes eventuate. For example, in its NDC system, Sweden includes automatic indexing for benefits at minimum pensionable age and for the increase in benefits for delayed retirement, but has not included automatic adjustment of the minimum pensionable age itself.¹⁹

In principle, any adverse effect on pension finances from increased life expectancy could be addressed by longer careers, since people could choose to spend part of their longer expected lives in continued work. That has not been the case in practice, for two reasons. First, retirement ages have generally decreased while life expectancy increased. Retirement is a normal good, that is, one for which demand increases as people’s incomes rise, so that increased demand for retirement has at least partially offset the effects of improved health and mortality. Second, as discussed earlier, even if people did extend their careers, the effect on pension cost would be limited in a system that is roughly actuarially fair.

If benefits are to be adjusted for mortality changes, automatic adjustment should be based on three principles.

- The rules should relate to date of birth not to the date of retirement, otherwise there will be a wave of retirement just before any reduction in the generosity of benefits. Such an incentive is inefficient.

¹⁹ Thus the endogenous variable is not the minimum pensionable age but the size of the pension. In a world of rationality this would not be an issue. However, a person whose personal discount rate exceeds the rate of actuarial adjustment of the pension will retire as soon as possible, creating potential pensioner poverty. Thus consideration needs to be given to adjusting the minimum pensionable age as well.

- Changes should be made annually, to avoid large changes in benefits across nearby cohorts. Large changes are inequitable and politically difficult, since benefits could otherwise differ significantly between people born in successive years, sometimes only days apart. The combination of large changes and rules determined by date of retirement would exacerbate the inefficient incentive to early retirement.
- Rules for changing benefits should be explicit, rather than adjusting the system in the light of experience. Greater predictability and decreased political pressures seem better with automatic adjustment with given rules. Moreover, automatic adjustments may function better – and politically more easily – if adjustments are based on actual mortality outcomes rather than projections. Nevertheless, there always remains the option of legislation to change whatever the automatic rules produce.

The legislated increase in women's pensionable age in the UK illustrates all three points.²⁰

One way of adjusting pension systems to longer lives after retirement is to reduce the average level of pensions. This is the approach taken by policy makers in Sweden, is inherent in a fully-funded defined-contribution system, such as in Chile, and was proposed for the US by a commission appointed by President Bush. Or it is possible to combine reduced benefits and increased contributions, as proposed for the US by Diamond and Orszag (2004, 2005).

As well as adjusting contributions and/or benefits, it is possible to increase the earliest pensionable age, one of the recommendations in the UK Pensions Commission (2005), discussed more fully in the paper by Hills (this journal). Since the State Pension Age is both the earliest entitlement age and the age for full benefits, such an increase is both a reduction in benefits for any given age of starting benefits and an increase in the minimum age at which a retirement benefit can be claimed. As discussed earlier, the latter adjustment raises complex issues, because the factors determining how many workers gain and how many lose vary with the increase both in life expectancy and in the level of earnings. A simple rule making minimum pensionable age proportional to life expectancy has advantages in terms of transparency but may be sub-optimal in theoretical terms: people are living longer, adding to the cost of pensions, but that effect is partially offset by the fact that people are better off than in the past, and so can afford to spend more on retirement. Given the absence of a clear rule relating optimal pensionable age to life expectancy, periodic adjustment of the minimum pensionable age – perhaps based on recommendations from an automatically-created cross-party commission – may be better than automatic adjustment.

²⁰ The reform was announced in 1991. The key date is 6 April 1950. For women born before that, State Pensionable Age will continue to be 60. Pensionable age for a woman born on 6 May 1950 (i.e. one month after the key date) would be 60 years and one month, for a woman born of 6 June 1950, 60 years and two months, and so on. Thus for women born on or after 6 April 1955 pensionable age will be 65.

3 Finance and funding

As discussed earlier (see note 8), finance and funding²¹ are areas of intense and wide-ranging controversy. This section concentrates on three areas we regard as salient: implicit and explicit debt; the relation between funding, saving and growth; and the proper way to compare returns to PAYG and funded systems.

3.1 Implicit and explicit debt

‘Implicit pension debt’ has become part of the vocabulary of international dialogue on pensions – unfortunately without a standard definition, making it a source of much confusion. The core argument, quite correctly, is that the pension promises of government have a future cost. However, there are multiple measures reflecting the long-term financial position of a pension system. There are three separate frameworks for considering the projected future of a pension scheme financed fully from dedicated revenues.²²

From the perspective of the finances of the scheme, one wants to look at the present discounted value (PDV) of expenditures minus revenues as a measure of whether current law is plausibly sustainable over a reasonable length of time. At issue within this frame is the choice of an appropriate time horizon. Traditionally in the US a 75-year horizon has been used. This is long enough to allow considerable time for adjustment if the finances are thought to be out of balance. Recently there has been a push for focusing on an infinite horizon. One of the arguments for doing so is that with a 75-year horizon, imbalances in annual flows over the end of the horizon and beyond mean that later calculations – using a 75-year horizon that then extends further into the future – will find imbalances even if the current picture shows balance. On the other hand, projecting beyond 75 years involves such uncertainty that it is not clear how much relevance the approach has for setting current policies. A useful and widely-used compromise is to preserve the 75-year horizon but add the condition that the projected finances not be deteriorating at the end of the horizon, referred to as sustainable solvency.

From the perspective of the treatment of future generations, the use of a year-by-year calculation for 75 years can be supplemented by a cohort-by-cohort calculation for all cohorts already part of the scheme. The former is referred to as an open-group measure, the latter as a closed group measure. Such a calculation is instructive about distribution, although evaluating any particular outcome requires consideration of the different positions of different cohorts in terms of earnings levels and life expectancies.

²¹ We use the term financing to refer to the cash flow used each year to pay benefits, and the term funding to indicate financial assets held by the pension system.

²² There are further complications for a scheme that also makes use of general revenues.

A third perspective is the cost of the system as a whole relative to the economy as a whole. That is, even if a scheme is sustainably financed from dedicated revenues, one might conclude that it is too large or too small. Consideration of size needs to reflect the needs of retirees relative both to their earlier earnings and to the positions of contemporaneous younger cohorts and the alternative uses of the resources being raised to finance the pensions.

Whatever frame is used, there are three natural measures of pension spending – pounds, percentage of payroll (or taxable payroll) and GDP. We think the direct measure in pounds is not a good way to communicate the position of a scheme – it is hard to distinguish among different very large numbers. The other two measures are appropriate for the different frames of the financing of the scheme itself and its role in the economy.

Since pension systems are to be relied on by workers, legislated changes should be infrequent and should have reasonable lead time. Hence the importance of projections and measures that make the public aware that a change of some type will be needed, even if cash flows are positive in the near term. A calculation of financial imbalance is thus important.

Whether referring to this as implicit debt is a useful way to communicate is a separate issue. The term is useful because it reminds people that explicit debt is not the only claim being made on future generations. However, it has led some analysts to treat implicit debt as fully (or nearly fully) equivalent to explicit debt. In particular, some analysts who favour individual accounts have argued that government should issue new debt equal to the amount of implicit debt, in order to place assets into individual accounts, viewing such a step as creating accounts with no cost. However, this argument is problematical. It is important to recognise that such an action has real economic effects. Government can lower implicit debt as part of a pension reform, and many governments have done so. In contrast, once explicit debt has been issued in place of the implicit debt, the value of the debt can be lowered only by actions, such as repudiation or increased inflation, that affect more than just the initial owners of the debt. Moreover, unless issued as consols,²³ for which the market is unclear, explicit debt needs to be rolled over repeatedly, creating vulnerability to bond market conditions. With a genuine difference in government options, it is not surprising that markets are likely to perceive explicit debt as different from implicit debt and hence to respond with some reluctance to a large and rapid increase in the supply of bonds, thus increasing the interest rates at which the government borrows.²⁴ In short, implicit and explicit debt are not equivalent.

²³ I.e. government bonds of infinite duration.

²⁴ A further difference arises once we consider matters in a more complex model than one with complete certainty. Uncertainty about future conditions in bond markets is far more important when there is considerably more debt outstanding, as would follow from converting implicit debt into explicit debt.

In assessing such measures, a central observation is that it is never required to fund implicit pension debt fully, just as a country does not ever need fully to pay off its explicit national debt. With explicit debt, what is critical is that the ratio of debt to GDP does not rise to the point where it induces large increases in interest rates and becomes unsustainable. Similarly, what is critical with implicit debt is that the contribution rate needed in the future does not rise so much that it undercuts participation to the point where the system becomes unsustainable. Giving some assets to a social security trust fund to reduce implicit debt may well be sensible, but allocating enough assets to bring the implicit debt down to zero is not a necessary condition for good policy.

The cost of pensions matters considerably. Excessive pension spending can reduce investment and cause major distortions which interfere with economic growth. Thus it is important to project future pension costs, but such projections must be interpreted correctly. Implicit debt is a useful concept but should not be given excessive weight.

3.2 Funding, saving and growth

Two sets of strategic arguments are frequently made for funded pensions: that funding increases growth (a macroeconomic argument); and that people regard their property rights as more secure if based on the ownership of financial assets (a political economy argument). For reasons of space, this section evaluates only the economic arguments.

In principle, funding increases growth if it increases national saving and/or improves the effectiveness of capital markets. A separate argument, also discussed below, is that funding assists adjustment to demographic change.

Funding and national saving

EFFECTS ON SAVING. An increase in national savings requires a decline in someone's consumption: raising contribution rates or cutting benefits now will lower the consumption of today's workers or today's pensioners, to make possible lower contribution rates or higher benefits in the future. In this case, increased funding raises the burden on current generations in order to lower the burden on future generations, similar to a budgetary decision to increase taxes or cut public spending so as to reduce public debt.

The process of building a fund may add little or much to national savings. There are two questions: does funding increase saving; and, if so, is the result welfare enhancing? On the first, the impact of an increase in funding on national savings can be anything from negative to large positive depending on the reaction of private savers and of the rest of the government budget. If workers are obliged to pay contributions into funded accounts, they

might respond by reducing their voluntary savings, in which case the introduction of a mandatory funded element has little effect on national savings. If shifting contributions from a central fund to individual accounts leads workers to save less because individual accounts bear a stronger parallel to voluntary savings, the impact can be negative. On the other hand, workers may continue with their voluntary saving, so that national savings increase. Indeed, workers doing no saving and having limited borrowing ability will sustain their (zero) private savings in the face of accounts financed by additional mandatory contributions. The overall outcome will depend on the balance of different responses by different groups of workers.

In addition to private responses, it is necessary also to consider government responses of changing taxes or spending in areas other than pensions. New revenues going into individual accounts may lead the government to spend more in other areas so that, again, there is little or no increase in national savings.²⁵ Transfers of revenue from a central fund to individual accounts may be financed by additional government borrowings so there will be no substantial increase in savings. Some argue that transferring revenues from a central fund to individual accounts will decrease spending elsewhere, since it may impact on reported deficit measures. While we have considerable evidence on individual savings responses to pension design, it is much harder to reach a conclusion that is well supported for the response of the rest of the government budget, a response, moreover, that is likely to vary with the ability of the government to do additional borrowing.

If funding does raise national savings, is the outcome beneficial? Specifically, does it make sense for an economy to raise contributions or reduce benefits now in order to have lower contributions or higher benefits in the future. Increased funding through lower benefits or higher contributions necessarily redistributes across generations. Thus, there can be no universal answer about whether funding raises welfare. Each country must consider the question in the context of its own circumstances and priorities, including its current saving rate and anticipated growth in earnings. If the saving rate is already high and growth rapid, it may make little sense to adopt a policy to increase savings even further.

The discussion thus far has concerned funding to increase savings. An alternative is for government to use workers' contributions to pay benefits, while simultaneously placing into workers' accounts a matching value of newly issued government debt. The effect on national savings of the latter policy is similar to that of a PAYG system. Thus it is important to distinguish between funding whose purpose is to increase savings (sometimes referred to as *broad funding*) and funding based on newly issued bonds – *narrow funding* – which does not increase savings.

²⁵ The impact on national savings can even be negative if workers reduce private savings to offset the mandate and yet government uses the easier availability of funds to decrease public savings.

ANALYTICAL ERRORS. Discussion of funding is prone to errors. Analysis often starts by comparing alternative steady-state outcomes with different degrees of funding, and follows with a comment that there is a transition cost to reach a steady state with a higher level of funding. This approach can be misleading. It gives little insight into the tradeoff between the benefits of having a larger fund and the cost of building one. Indeed the term ‘transition cost’ suggests something small, even for a transition period of decades (as in Chile, with more decades to come).²⁶ It is more informative to analyse funding as above, by considering the implications of increasing funding today to have some combination of lower taxes or higher benefits in the future. This way of posing the issue does not focus on funding per se but on the tax, benefit, and debt decisions which should underpin any decision to increase funding.

A separate, widely-made, but incorrect case for funded accounts is the argument that a funded system is better if the rate of interest (i.e. the return on pension funds) exceeds the rate of growth (i.e. the return to a PAYG system). Once the analysis is done fully (section 3.3), there is no gain for everyone from funding per se, but an intergenerational redistribution. That is, the comparison is basically the same incomplete analysis of considering only the steady state without considering the adjustment to a new steady state.

Another suspect argument is that ‘a multipillar structure allows for tactical sequencing, strategic bundling, packaging, and compensation and thus is useful for overcoming resistance to reform’ (Holzmann and Hinz, 2005, p. 42). The argument suggests that a gain for everyone is available by bundling analysis of static efficiency improvements in pension rules with funding. But those efficiency gains are available without funding, by better design of a PAYG system, so this is misattribution, unless it really is the case that the efficiency gains are not available politically without such bundling. In politics bundling does matter, but may not be necessary for reform, and indeed may make reform more difficult (see, for example, Boersch-Supan, 2006, on the German case). In any case, political arguments are separate from economic arguments and should be clearly labelled as such.

Finally, some analysis implies that funding necessarily requires individual accounts. Again, this is mistaken. If policy makers want more funding, there are many ways to bring it about. Sweden has funded a diversified central portfolio within a DB system for years, and similar arrangements have been started in Canada and Switzerland. Funding a central fund within a DC system has been done by the provident funds of Malaysia and Singapore. And funding is possible also with worker choice over portfolios offered by private providers, as pioneered in Chile and followed in other Latin American countries, implemented in the UK, and recently in Sweden. Thus, choosing a level of funding and a degree of portfolio diversification is economically unrelated to the choice between DB and DC systems or between individual and more broadly-based accumulations.

²⁶ See the discussion of Chile in the paper by Arenas and Mesa Lago, this journal.

Funding and capital market development

Alongside any effects on saving, funding assists growth if it helps to improve the efficiency with which savings are channelled into investment.

In two polar cases the argument clearly fails. In advanced countries, financial markets are highly developed, so that mandatory pension savings are unlikely to bring about substantial further improvements.²⁷ Contrariwise, in countries with very limited institutional capacity, the existing financial infrastructure is too weak to risk the pensions of large numbers of workers by mandating funded individual accounts.

Between the two is a range of country capacities where there is the potential to improve capital markets but also the risk, without enough improvement, that workers will not get good returns on their contributions, or that government will have to bear the cost of bailing out the pension system. The risk is easy to comprehend: inadequate markets can yield low returns. They also have much higher costs than better-developed ones, a point of particular relevance to small accounts. Adequate markets require significant government regulation. Indeed it is argued that the US has the best-functioning capital markets because the US has the best-regulated capital markets. It also helps to have a large market, by being based in a large economy. As discussed shortly, poor markets not only hurt pensioners but may also imply a worse allocation of investment than would occur with less formal ways of allocating savings to investment. The possibility of gain is also easy to comprehend, since better-functioning capital markets increase economic efficiency and so economic growth. What is critical to the possibility of gain is a sustained effort to improve the regulation of markets and the functioning of the economy generally, an effort that may be enhanced by committing the funds of workers and so their political interests.

An alternative approach is to encourage voluntary pensions as a stimulus to market development, particularly where the economy is large enough that voluntary savings can reap economies of scale.

Exploring the issue in more detail, there are four channels through which savings are allocated to investment:

- (a) Market transactions – purchasing newly issued bonds and stocks (as opposed to trading existing ones).
- (b) Intermediation – deposits in banks and other intermediaries, which are then lent to investors.

²⁷ Insofar as the set of investors making good use of capital markets is increased by such funding, there is a potential to spread risks more widely and so more efficiently.

(c) Direct pairwise loans: to people starting small businesses, from friends and families; to people one transacts with through trade credit; or seller-provided credit more generally.

(d) Saving to finance ones own investment.

In a country where market structure is weak and the banking system functions poorly, progress is not likely to come through (a) or (b). Individual accounts, by taking more money from workers, cut down pairwise loans and own investment, with no places where pension funds could allocate savings that offer any evidence of functioning better in the short run.

A second set of arguments about the role of pension funds in improving the efficiency of capital markets is whether transferring shares of newly-privatised state-owned enterprises to the social security trust fund (as in some of the former-communist countries in Central and Eastern Europe) improves corporate governance – a key ingredient in economic efficiency and economic growth in market economies.²⁸ High-quality governance needs good legislation, effective oversight by the regulatory authorities and effective oversight and exercise of voting rights by share owners. In advanced countries, centralised pension funds (e. g. defined benefit or provident funds) can result in additional players in corporate governance which can influence outcomes, for better or worse.²⁹

Funding and demographic change

Though the point was shown to be flawed many years ago,³⁰ the argument that funding necessarily assists pension finance significantly in the face of demographic change still reappears. The argument needs to be considered in the context of a decline both in fertility and elderly mortality.

Suppose that a large workforce is followed by a smaller workforce. In a pure PAYG scheme the revenue from a given social security contribution rate falls, creating upward pressure on the contribution rate, downward pressure on the level of pensions, or both. This problem is well-understood and not controversial.

It is argued that funding can ease the problem: each member of the large workforce in period 1 builds up pension savings; the DC pension available for a representative worker is exactly what can be covered by those savings; if there is a large number of such workers, this is not a problem, it is argued, because each worker accumulates enough, on average, to pay

²⁸ See the symposium on corporate governance in *Oxford Review of Economics Policy*, 21/2, Summer 2005.

²⁹ The California Public Employees' Retirement System (CalPERS), for example, with over \$150 billion in assets, takes an active shareholder interest (see Boersch-Supan and Winter, 2001).

³⁰ Barr (1979); for a recent restatement see Barr (2000, IIA).

for his or her own pension. This argument is correct in terms of *finance* but may fail to provide workers with the *consumption* they expect in old age. With PAYG, the shortfall comes through a decline in social security contributions. With funding the mechanism is less direct but has the same cause: unless a decline in the number of workers has no effect on output, output will fall; and if output falls consumption and/or investment must fall. Lower rates of return or higher prices deny pensioners the consumption they expected; or mandatory increases in pension savings by workers reduce their consumption by more than they would choose; or the increase in the combined consumption of workers and pensioners is at the expense of investment, and hence puts future growth at risk. As noted earlier, PAYG and funding are both mechanisms for organising claims on future output; since demographic change generally affects that output, it generally causes problems for pension schemes however they are organised.

An even closer parallel exists if the birth rate is stable, but the life expectancy of pensioners increases. The effect is to increase the number of pensioners per worker. With pure PAYG this increase requires a higher contribution rate or lower monthly benefits to maintain the balance of the scheme. With funding and no change in interest rates, the sustainable level of monthly benefits is lower if the retirement lifetime is longer. With an interest rate that exceeds the return implicit in the PAYG system the precise effects may differ somewhat, but the character of the problem is the same in both cases.

What matters is not financial accumulation but output. If output increases, it becomes easier to meet the claims both of workers and pensioners. The solution to population ageing lies not in funding per se but in output growth.

Conclusion

The relationship between funding and growth is neither simple nor automatic.

1) Does funding increase saving? As discussed, it will fail to do so if an increase in mandatory pension saving is offset by a decline in voluntary saving or a decline in the saving of government elsewhere in the budget. Thus saving may or may not increase – the amount of increased national savings has a complex relationship with the amount of increased funding.

2) How much will an increase in saving increase output? The simplest argument is that a move to funding (a) increases savings, which (b) increases investment, which (c) increases output by the marginal product of capital. These links hold in many circumstances, but not always or necessarily and not with a simple connection.

- As noted, a move to funding does not necessarily increase saving.

- The link between an increase in saving and increased investment is complex – some savings will simply increase prices of existing assets. In the 1970s a British trade union famously invested part of its pension fund in valuable paintings, creating no new factories or machinery. Part of increased saving can drive up the prices of assets in limited supply, such as urban land.
- An increase in investment may not increase output by much. Inefficiencies in capital markets may make the marginal product of investment low, as in the communist countries in Central and Eastern Europe and the former Soviet Union, which all had rates of investment that were exceptionally high by Western standards, yet had growth that stagnated, and in some countries turned negative. Moreover, returns on financial assets include an adjustment for bearing risk. Thus it is wrong to use financial returns as a measure of the return to society without incorporating an adjustment for risk, and hence wrong to use the expected return on stocks as the gain from increased funding.

3) Is such a policy optimal? The fact that an increase in funding may increase output, does not mean that the policy is necessary welfare improving. Beneficial effects will depend on country-specific features.

- Are savings in the country lower than optimal and/or are pension funds likely to improve the allocation of savings or corporate governance?
- Intergenerational issues are another aspect of the previous point. Increased investment through increased funding implies lower consumption in the present. Thus funding that increases growth is worthwhile only if the fall in current consumption is less valuable than the increase in future consumption.

4) Is such a policy feasible? The answer will depend in part on whether the country has the necessary institutions? Does it have the necessary skills in allocating pension funds, skills in administering pension accounts, and the capacity to regulate financial markets?

In sum, the argument is not that funding is bad policy, nor that it cannot help with population ageing, but that its helpfulness is contingent on beneficial effects on growth and on country-specific factors. Funding may be important for economic growth – but the case has to be made in each country, not just assumed or asserted.

3.3 Comparing the returns to PAYG and funding

Some analysts compare the long-run return on assets with the rate of growth, which is the long-run return in a PAYG system.

‘In contrast to the 2.6-percent equilibrium return on Social Security contributions, the real pretax return on nonfinancial corporate capital averaged 9.3 percent over the

same ... period.... [As a result], forcing individuals to use the unfunded system dramatically increases their cost of buying retirement income' (Feldstein, 1996, p. 3).

Since long-run rates of return exceed growth rates the higher stock market return is sometimes presented as a pure gain. This argument is flawed because it does not compare like with like. A fuller analysis considers (a) the costs of the transition from PAYG to funding, (b) the relative risks of the two systems, and (c) their respective administrative costs.

Inappropriate comparison of steady states

If proper account is taken of the costs of transition from a PAYG to a fully funded scheme there is generally an equivalence between the rates of return in the two schemes.³¹

The flaws in the argument that pensioners are better-off under funding if the stock-market return exceeds real wage growth can be seen clearly if policy makers are considering establishing a pension system in a brand new world. In a funded system, in contrast with PAYG, the first generation of pensioners will not receive a pension; thus it is mistaken to present the gain to pensioners after the first generation as a Pareto improvement. Mostly, however, what is being discussed is a move from an existing PAYG scheme towards funding. In that case, including the transition costs of the change gives the same picture – some are helped and some are hurt. A central question is where those costs will fall.

Table 1: A Simplified Pay-As-You-Go System

Period	Generation			
	A	B	C	D
1	+\$1	-\$1		
2		+\$1	-\$1	
3			+\$1	-\$1
4				+\$1

Source: Orszag (1999, p. 9)

Case 1: constant benefit rules; transition costs financed by public borrowing: in Table 1, each generation pays \$1 in contributions when young and receives \$1 in pension when old. In period 1, the \$1 pension of older generation A is paid by the \$1 contribution of younger

³¹ The argument draws on Orszag (1999), a non-technical summary of results originally established by Breyer (1989), and applied to the USA by Geanakoplos, Mitchell, and Zeldes (1999), and Belan and Pestieau (1999).

generation B. In period 2, when generation B is old, its pension is paid by the contributions of young generation C. Now suppose that the real rate of return on assets, i , is 10 per cent, and imagine that we are generation C. Under a PAYG scheme we pay \$1 in contribution in period 2 and receive \$1 pension in period 3; the real rate of return is zero. In contrast, with an individual account we save \$1 in period 2 and get back \$1.10 in period 3; the real rate of return, it appears, is 10 per cent.

The flaw in the argument is that if generation C contributes to its own funded accounts, generation B's pension must be paid from some other source. If that source is government borrowing, generation C receives a pension of \$1.10 but has to pay interest of 10 cents on the borrowing which financed generation B's pension. The real return – as under PAYG – is zero. The lower return on the PAYG system is not the result of some inherent flaw, but is precisely the cost of the initial gift to generation A. Formally (see Breyer 1989, Belan and Pestieau 1999), there is an equivalence between the two schemes if the move to funding is considered not in isolation but alongside the cost of financing the change. Thus generation C is not made better-off by a move to individual accounts.

‘[F]alling money’s worth in this model is *not* due to the aging of baby boomers, increased life expectancy, or massive administrative inefficiency, but rather to the simple arithmetic of the pay-as-you-go system’ (Geanakoplos, Mitchell and Zeldes, 1999, p. 86, emphasis in original).

Case 2: constant benefit rules; transition costs financed by taxation: suppose that we are generation C: in period 2 we put our contribution of \$1 into an individual funded account, and the \$1 pension of generation B is paid in part from a tax on generation C. The pension we receive as generation C is \$1.10. But generation C paid part of generation B's pension through a tax, and so has less wealth to finance retirement. The real return on the assets is 10 per cent but this does not imply a return of 10 per cent on the combination of the mandatory savings and the tax paid toward generation B's pensions.

Case 3: no benefits to the transition generation: another way to finance the transition is to throw generation B out of the lifeboat by not paying their pension at all. Generation C and onwards enjoy a 10 per cent real return, but those gains are at the expense of generation B, on whom the entire cost of transition is concentrated. In this case, the cost of the gift to generation A is offset by the negative gift to generation B.

The fundamental point is that there is a zero-sum game between the first generation and subsequent generations. The cost of the gift to the first generation can be placed entirely on the transition generation of pensioners (generation B) by reneging on PAYG promises; or entirely on the generation of workers at the time of transition (generation C) by financing generation B's pension out of taxation; or spread over succeeding generations by financing

the transition through borrowing. It is possible to alter the time path of the cost, but not its total. Again, the only way out of the impasse is if a move towards funding leads causally to higher rates of growth, an issue on which, as discussed earlier, controversy continues.³²

Adjusting for differences in risk

The cost of financing the transition is only part of the comparison between PAYG and funding. A second element is risk, the key point being that the real return to PAYG and to funded schemes should be adjusted downwards to account for risk. The simplest approach is to assume a bond rate of return on stocks, although this understates the value to those with no outside assets.

Alongside market risks are political risks, which are present with both kinds of systems. It is true that government action can increase risk through various forms of government failure. Equally, however, some adjustments by governments are precisely to address, at least partially, the risks that individuals face. Formal analysis of such political risks and how they differ across systems has not advanced very far.

Controlling for administrative costs

The evidence that the administrative costs of individual accounts are higher – often considerably higher – than PAYG schemes is well-established.³³ The importance of administrative costs should not be underestimated. Under plausible assumptions, over a working life, an annual administrative charge of 1 per cent of a person's pension accumulation will reduce the total accumulation by about 20 per cent.³⁴

Conclusion

In assessing proposals for pension reform a central analytical point is to be clear what question is being asked. Feldstein (2005), for example, argues that social security in the USA reduces savings. The analysis makes a steady state comparison, that is, compares the economic situation in the USA today with what it would have been in an alternative steady state with funded pensions. Thus the underlying question is: how does welfare in steady state B differ from that in steady state A? Most of the analysis in this paper is about a different question: what are the welfare effects of moving from steady state A to steady state B? Either question (and answer) is legitimate. What is not legitimate is to apply the answer from one question to the other.

³² The analysis in the text has considered a simple setting, in particular without taxes on the return to capital. Such taxes complicate the analysis but do not alter the need to consider winners and losers in any reform.

³³ For UK data, see UK Pensions Commission (2004, Table 6.9 and surrounding discussion).

³⁴ For details, see Barr and Diamond (forthcoming, Ch. 3).

The conclusion is not that a move to funding is always bad policy, but that its desirability cannot be established by simple comparison of rates of return.

- A move from PAYG towards funding should take proper account of the costs – both their total and their distribution – of transferring from one steady state to another, of differences in risk, and of any discrepancy in administrative costs.
- All three adjustments remain relevant to the choice of pension regime in a hypothetical new country where the issue of transition costs is replaced by the issue of whether to give benefits to people who have already retired or are close to doing so.

Atkinson (1999, p. 8) points out that critics of the welfare state tend to consider its costs without taking account of its benefits.

‘The emphasis by economists on the negative economic effects of the welfare state can be attributed to the theoretical framework adopted ..., which remains rooted in a model of perfectly competitive and perfectly clearing markets. [This] theoretical framework incorporates none of the contingencies for which the welfare state exists... The whole purpose of welfare state provision is missing from the theoretical model’.

The point here is precisely similar: that the benefits from a move to funding should not be considered in isolation but alongside the relevant costs.

4 Distributional issues

4.1 Sharing risks

The different pension arrangements discussed in section 1 share risks differently.

In a pure defined-contribution scheme, a person’s pension, given life expectancy, etc., is determined by the size of her lifetime pension accumulation. Thus the individual faces all the risks discussed earlier: macroeconomic shocks, demographic shocks, political risks, management risk, investment risk and annuities market risk.

Defined-benefit schemes share risks more broadly. A person’s pension may, de facto, be wage indexed until retirement. The employee contribution is generally a fraction of his/her salary. Thus the risk of varying rates of return to pension assets falls on the employer, and hence on some combination of the industry’s current workers (through effects on wage rates), its shareholders and the taxpayer (through effects on profits), its customers (through effects on prices) and/or its past or future workers, if the company uses surpluses from some periods to boost pensions in others.

With social insurance, risk is shared yet more broadly. The costs of adverse outcomes can be borne by the pensioner through lower monthly pensions or a shorter period of

retirement, by contributors through higher contributions, or by future pensioners or contributors if the pension fund is able to borrow.

Where pensions are at least partly tax financed, the risk falls on pensioners, if monthly pensions or retirement duration are reduced, or on contributors, or on current taxpayers or, via government borrowing, on future taxpayers.

In the latter two cases, the institutional structure is explicitly intended to spread risk, not an afterthought for poor outcomes.

4.2 Sharing burdens across generations

Since consumption plus investment adds up to national income, a policy which raises national savings will lower consumption. Conversely, a policy to give benefits to current retirees is designed to raise their consumption and so decreases investment, aggregate output held constant. Thus decreased savings is a necessary implication of higher pensions, not a surprise unlucky outcome. If the additional resources generated by any increase in savings are used to raise future benefits or to reduce future contributions, the effect is to increase future consumption. Such a policy thus redistributes across cohorts. To evaluate whether such redistribution is worthwhile, one needs to consider not only the return on the additional capital, but also the extent to which consumption would have grown anyway in the absence of the policy to increase savings. One also needs to consider the distribution within each cohort of the decreases in consumption now and the increases in the future.

More specifically, consider increasing the contribution rate now so as to have a lower contribution rate later. Workers, paying a higher contribution rate, have lower consumption. Later workers, paying a lower contribution rate, have higher consumption. How should we evaluate this redistribution from today's workers to later cohorts of workers? There are three parts to the comparison.

- How much does each worker value the consumption change? With diminishing marginal utility, the higher the level of consumption, the lower the value of an increase. Thus with rapid growth, moving consumption into the future is less valuable than with lower growth. Similarly, a country with a high savings rate has low consumption rate relative to its income. If the savings rate is lower in the future, as the country moves to a higher living standard, then again the increase in consumption lowers the gain from delaying consumption.
- The greater the investment needs of the country, the greater the marginal product of capital is likely to be, assuming adequate markets for allocating investment. But a country with a high savings rate, and so a high investment rate, is likely already to be investing in opportunities with the highest rates of return. Thus the return to yet higher rates of saving may not be so high.

- The pure public weighting of different cohorts: it is common to discount across cohorts as well as, separately, discounting for consumption during a worker's life.

Thus, increasing national savings by raising contributions or lowering benefits redistributes toward future cohorts. These distributional effects should not be ignored in evaluating proposals for pension reform.

5 Conclusion

If nothing else, we have probably succeeded in convincing the reader that analysis of pension reform proposals is complex and that overly simple incomplete comparisons are not a sound way to proceed.

The central arguments are:

- Pensions have multiple purposes. They may be given different weights, but policy needs to bear them all in mind.
- Formulating pensions policy in a first-best framework is simple, but is a bad guide to policy design in a world with major market imperfections. The analysis of labour-market effects of pensions in section 2 is thus framed in a second-best context.
- A move from PAYG towards funding may or may not be welfare-improving, depending on a series of country specifics. The analysis of any such move needs to take account of the costs of moving from one steady state to another, rather than simply comparing two steady states. It also needs to take account of differences in risk and of the significant differences in administrative costs of different types of pension arrangement. And it needs to take account of distributional effects within and across generations, including the distribution of risks.

References

- Aaron, Henry J. (1966), 'The Social Insurance Paradox', *Canadian Journal of Economics and Political Science*, XXXII (3), August, 371-4; repr. in N. Barr (ed.) (2001) *Economic Theory and the Welfare State*, Edward Elgar Library in Critical Writings in Economics, Vol. II, 79-82.
- Atkinson, A. B., 1999, *The Economic Consequences of Rolling Back the Welfare State* (Cambridge, Massachusetts: MIT Press).
- Barr, Nicholas (1979), 'Myths My Grandpa Taught Me,' *Three Banks Review*, No. 124 (December), pp. 27–55, reprinted in Nicholas Barr (ed.) *Economic Theory and the Welfare State*, Edward Elgar Library in Critical Writings in Economics, (Edward Elgar, 2001), Vol. II, pp. 83-111.
- Barr, Nicholas (2000), 'Reforming Pensions: Myths, Truths, and Policy Choices', Working Paper WP/00/139, Washington DC: International Monetary Fund, 2000, downloadable from <http://www.imf.org/external/pubs/ft/wp/2000/wp00139.pdf>.
- Barr, Nicholas (2001), *The Welfare State as Piggy Bank: Information, risk, uncertainty and the role of the State*, Oxford and New York: Oxford University Press.
- Barr, Nicholas and Diamond, Peter (forthcoming), *Reforming Pensions: Principles and an Application to China*, New York and Oxford: Oxford University Press.
- Barr, Nicholas and Rutkowski, Michal (2005), 'Pensions', in Nicholas Barr (ed.) (2005), *Labor Markets and Social Policy: The Accession and Beyond*, Washington DC: The World Bank, pp. 135-170.
- Belan, Pascal, and Pierre Pestieau (1999), 'Privatizing Social Security: A Critical Assessment,' *The Geneva Papers on Risk and Insurance—Issues and Practice*, Vol. 24, pp. 114–130.
- Boersch-Supan, Axel H. (2006), 'The German public pension system: How it will become and NDC system look-alike', in Holzmann, Robert and Palmer, Edward (eds), *Pension Reform: Issues and Prospects for Non-Financial Defined Contribution (NDC) Schemes*, Washington DC: World Bank, pp. 573-610.
- Boersch-Supan, Axel H., and Winter, Joachim K. (2001), Population Aging, Savings Behavior and Capital Markets, NBER Working Paper No. 8561, October.
- Breyer, Friedrich (1989), 'On the Intergenerational Pareto Efficiency of Pay-As-You-Go Financed Pension Systems,' *Journal of Institutional and Theoretical Economics*, Vol. 145 (December), pp. 643–58.
- Burtless, Gary (2002), 'Social Security Privatization and Financial Market Risk', in Ihori, T., and Tachibanaki, T., *Social Security Reform in Advanced Countries*, London and New York: Routledge.

Diamond, Peter A. (2002), *Social Security Reform*, Oxford and New York: Oxford University Press.

Diamond, Peter A. (2004), 'Social security', *American Economic Review*, 94/1 (March), 1-24.

Diamond, Peter A. and Orszag, Peter R. (2004), *Saving Social Security: A balanced approach*, Washington DC: The Brookings Institution.

Diamond, Peter A. and Orszag, Peter R. (2005), 'Saving social security', *Journal of Economic Perspectives*, 19/2, Spring, pp. 11-32.

Feldstein, Martin S. (1996), 'The Missing Piece in Policy Analysis: Social Security Reform,' *American Economic Review*, Vol. 86 (May), pp. 1-14.

Feldstein, Martin S. (2005), 'Rethinking social insurance', *American Economic Review*, 95/1, March, pp. 1-24.

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes (1999), 'Social Security Money's Worth,' in *Prospects for Social Security Reform*, ed. by Olivia S. Mitchell, Robert J. Myers, and Howard Young (Philadelphia: University of Pennsylvania Press).

Holzmann, Robert and Hinz, Richard (2005), *Old Age Income Support in the 21st Century: An international perspective on pension systems and reform*, Washington DC: The World Bank.

New, Bill (1999). 'Paternalism and Public Policy', *Economics and Philosophy*, 15: 63-83.

Orszag, Peter, 1999, *Individual Accounts and Social Security: Does Social Security Really Provide a Lower Rate of Return?* (Washington: Center on Budget and Policy Priorities), downloadable from <http://www.cbpp.org>.

Orszag, Peter R., and Stiglitz, Joseph E. (2001), 'Rethinking Pension Reform: 10 Myths About Social Security Systems', in Holzmann, Robert and Stiglitz, Joseph E., with Fox, Louise, James, Estelle, and Orszag, Peter R. (eds), *New Ideas About Old Age Security: Toward Sustainable Pension Systems in the 21st Century*, Washington DC: The World Bank, pp. 17-62.

Palmer, Edward (2006), 'What is NDC?', in Holzmann, Robert, and Edward Palmer, eds. 2005. *Pension Reform through NDCs: Issues and Prospects for Non-Financial Defined Contribution Schemes*. Washington, DC: World Bank, pp. 17-34.

Samuelson, Paul A. (1954), 'An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money', *Journal of Political Economy*, 66/6 (Dec.), 467-82; repr. in N. Barr (ed.) (2001) *Economic Theory and the Welfare State*, Edward Elgar Library in Critical Writings in Economics, Vol. II, 63-78.

UK Pensions Commission (2004), *Pensions: Challenges and Choices: The First Report of the Pensions Commission*, London: TSO, downloadable from <http://www.pensionscommission.org.uk/publications/2004/annrep/index.asp>

UK Pensions Commission (2005), *A New Pension Settlement for the Twenty-First Century: Second Report of the Pensions Commission*, London: TSO, downloadable from <http://www.dwp.gov.uk/publications/dwp/2005/pensionscommreport/main-report.pdf>