



Københavns Universitet



Individual Savings Accounts and the Life-Cycle Approach to Social Insurance

Sørensen, Peter Birch; Hansen, Martin Ino; Bovenberg, A. Lans

Publication date:
2006

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
Sørensen, P. B., Hansen, M. I., & Bovenberg, A. L. (2006). Individual Savings Accounts and the Life-Cycle Approach to Social Insurance. Economic Policy Research Unit. Department of Economics, University of Copenhagen.

EPRU Working Paper Series

**Economic Policy Research Unit
Department of Economics
University of Copenhagen
Studiestræde 6
DK-1455 Copenhagen K
DENMARK
Tel: (+45) 3532 4411
Fax: (+45) 3532 4444
Web: <http://www.econ.ku.dk/epru/>**

Individual Savings Accounts and the Life-Cycle Approach to Social Insurance

**Peter Birch Sørensen, Martin Ino Hansen,
and A. Lans Bovenberg**

2006-03

ISSN 0908-7745

**The activities of EPRU are financed by a grant from
The National Research Foundation**

March 2006

INDIVIDUAL SAVINGS ACCOUNTS AND THE LIFE-CYCLE APPROACH TO SOCIAL INSURANCE

Peter Birch Sørensen, University of Copenhagen,
The Danish Economic Council, EPRU and CESifo*

Martin Ino Hansen, Statistics Denmark**

A. Lans Bovenberg, CentER, Tilburg University,
Netspar, CEPR and CESifo***

Abstract: Using Danish data, we find that about three fourths of the taxes levied to finance public transfers actually finance benefits that do not redistribute between people but redistribute income over the life cycle of individual taxpayers. This provides a rationale for financing part of social insurance via mandatory individual savings accounts. An account system that offers liquidity insurance and a lifetime income guarantee helps to alleviate the dilemma between insurance and incentives. To illustrate this, we analyse a specific proposal for reform of the Danish system of social insurance, involving the use of individual accounts. We estimate how the reform would affect the distribution of lifetime incomes, the public budget, and economic efficiency. Our analysis suggests that, even with conservative assumptions regarding labor supply elasticities, the proposed reform would generate a Pareto improvement and would imply only a minor increase in the inequality of lifetime income distribution.

Keywords: Social insurance, individual accounts, lifetime income distribution

JEL Code: H53, H55

*) Department of Economics, University of Copenhagen, Studiestraede 6,
1455 Copenhagen K, Denmark. E-mail: Peter.Birch.Sorensen@econ.ku.dk

**) Statistics Denmark, Sejrøgade 11, 2100 Copenhagen Ø, Denmark.

E-mail: mih@dst.dk

***) CentER and Netspar, Tilburg University,

P.O. Box 90153, 5000 LE Tilburg, The Netherlands. E-mail: A.L.Bovenberg@uvt.nl

INDIVIDUAL SAVINGS ACCOUNTS AND THE LIFE-CYCLE APPROACH TO SOCIAL INSURANCE

Peter Birch Sørensen, Martin Ino Hansen and A. Lans Bovenberg¹

1. Introduction

Many welfare states are under strain as a result of globalization, aging, and technological change biased against low-skilled workers. While these trends have received much attention, it is less frequently recognized that the changing nature of social risks also puts the welfare state under pressure. To illustrate, as the economy shifts from blue-collar work in industrial sectors to white-collar work in service sectors and knowledge-intensive activities, mental causes of sickness and disability become more prominent. These types of sickness and disability are less easy to diagnose and verify than those with physical causes, thereby increasing the danger of moral hazard. Moreover, changes in technology and in the organization of work have made many segments of the labour market more 'fluid', as people move more frequently between employers and as they enter and exit the labour force more often. In such a transitional labor market it becomes more difficult to verify whether a person is voluntarily or involuntarily out of work, again exacerbating the problem of moral hazard in social insurance. Thus, whereas the dynamic world economy confronts many people with increasing economic risks, the ability of the welfare state to offer security is weakened, as globalization increases the mobility of tax bases and as the changing nature of human-capital risks raises the costs of insuring these risks. The age-old trade-off between equity and efficiency as well as the related dilemma between insurance and incentives are more relevant than ever before.

Against this background, we analyze the merits of mandatory individual saving accounts that are supplemented by public liquidity insurance and public lifetime income insurance. Mandatory payments in personal saving accounts that finance social insurance payments (including deductibles and insurance premiums) replace taxes that are currently financing social-insurance benefits. At retirement, the remaining balances in

¹We are grateful to Steen Jørgensen and Anne Kristine Høj for methodological advice regarding the estimation of lifetime incomes. Any remaining shortcomings are our own responsibility.

the accounts are converted into an annuity, which is added to the ordinary public retirement benefit. If the account balance is negative at that time, the account is set to zero and the account holder simply receives the ordinary retirement benefit. We explore to what extent such accounts can improve incentives without substantially increasing lifetime inequality and lifetime risk. The motivation behind these accounts is that they facilitate consumption smoothing throughout the life cycle without creating substantial disincentives to work. These accounts therefore limit the inescapable labour-market distortions that are associated with lifetime redistribution, lifetime insurance and liquidity insurance. In particular, the accounts establish an efficiency-enhancing actuarial link between contributions and benefits for high-income and middle-income workers (who currently pay distortionary taxes partly to finance distortionary social benefits to themselves) without harming low-income workers who remain protected by the lifetime income guarantee. This actuarial link reduces the tax character of social security contributions and thus enhances incentives to work. The accounts also help to improve the trade-off between insurance and incentives by facilitating self-insurance over the life course. In particular, they allow people to shift the payment of deductibles in social insurance to the periods in which these costs can be more easily afforded. In this way, individuals can self-insure themselves over their life course – and thus do not have to rely on insurance that gives rise to moral hazard.

To motivate the life-cycle approach to social insurance adopted in this paper, we start by documenting the rather low degree to which current social-insurance programs redistribute lifetime incomes. Using Danish data, we find that about three-fourths of the taxes levied to fund public transfers merely finance benefits that redistribute income over the same taxpayer's life cycle rather than between different people. We also identify those social-insurance programs that imply a high degree of lifetime income redistribution and therefore would not be suited for financing via individual saving accounts. Following this empirical analysis, we explain the theoretical rationale for mandatory saving accounts for social insurance. We argue that a properly designed account system has the potential to generate a Pareto improvement, from not only an ex-ante perspective (i.e. before people know which shocks they are going to encounter during their lives) but also from an ex-post point of view (i.e. after all shocks have materialized). To illustrate how an

account system might work in practice, we proceed to lay out a specific proposal for a reform of the Danish system of social insurance, involving the use of individual accounts. We then estimate how this reform would affect the distribution of lifetime incomes, the public budget and economic efficiency. Our analysis suggests that, even with conservative assumptions regarding labor-supply elasticities, the proposed reform would indeed imply a Pareto improvement and would involve only a minor increase in the inequality of lifetime income distribution.

Other studies have explored lifetime distribution and how it is affected by public policy. Nelissen (1998) employs a micro-simulation model to investigate how social insurance affects income distribution on a lifetime basis in the Netherlands. Similar studies for the United States include Coronado et al. (2000), Gustman and Steinmeier (2000), and Liebman (2001). These studies find that social insurance helps to redistribute from the lifetime rich to the lifetime poor. The contribution of social insurance to the redistribution of lifetime incomes, however, is considerably smaller than what is suggested by its contribution to the redistribution of annual incomes. To illustrate, Nelissen (1998) finds that the Dutch social-insurance system reduces annual income inequality by 45 percent but lifetime inequality only by between 15 and 30 percent (depending on the discount rate and on the historical cohort considered). Ter Rele (2003) covers a broader range of Dutch governmental programs, but does not account for all heterogeneity and restricts himself to an analysis of the life cycles of six educational groups. He finds that public programs substantially reduce inequality on a lifetime basis. Whereas lifetime tax liabilities are proportional to lifetime incomes, lifetime benefits of public spending do not vary with lifetime income.

To estimate the degree to which the public sector redistributes lifetime income, this paper uses a method that was previously applied by Hussénius and Selén (1994), Falkingham and Harding (1996), O'Donoghue (2001), and Pettersson and Pettersson (2003). Section 2.3 explains this method and compares the results from these previous studies to our own findings.

Fölster (1997, 1999), Orszag and Snower (1997, 2002), Feldstein and Altman (1998), Orszag et al. (1999), Fölster et al. (2002), Stiglitz and Yun (2002), Sørensen (2003) and Bovenberg and Sørensen (2004) analyze the merits of various types of individual

saving accounts. Some of these studies investigate how individual accounts for the financing of unemployment benefits could improve labour-market incentives compared to a tax-financed system of unemployment insurance. Fölster (1997, 1999) estimates how individual accounts financing a broader set of social-insurance programs would affect the distribution of lifetime incomes in Sweden, given that the account system includes lifetime income insurance as well as liquidity insurance. We conduct a similar exercise for Denmark in this paper. In contrast to Fölster (1997, 1999), we also explore how the system of individual accounts affects work incentives, allowing us to compute the efficiency gains and the consequences for the government budget. Moreover, we design our proposal in such a way that it is Pareto improving.

The rest of the paper is structured as follows. Section 2 estimates how much the various social-insurance benefits redistribute lifetime incomes as compared to annual incomes. Section 3 investigates the advantages and disadvantages of individual accounts with a lifetime-income guarantee as an instrument to enhance labor-market incentives without harming the lifetime poor. Section 4 analyzes how individual accounts affect the distribution of lifetime incomes, the labor market, the public budget and economic efficiency. Section 5 concludes. Two technical appendices document some of the results reported in the main text.

2. Redistribution over the life cycle versus redistribution of lifetime incomes: a case study of the Danish welfare state

In the modern Western European welfare state a substantial part of the taxes paid by the average taxpayer finance public transfers that are channelled back to the same taxpayer at some point in the life cycle. This section presents a case study of the Danish welfare state indicating that only about one-fourth of social-insurance transfers helps to redistribute income from the lifetime rich to the lifetime poor. We also present estimates of the different degrees to which the various social-insurance benefits redistribute lifetime incomes as compared to annual incomes. Several benefits that appear to be highly redistributive in a cross-section analysis based on annual incomes turn out to have little redistributive power when we adopt a life-cycle perspective.

The analysis in this section provides the empirical motivation for the life-cycle ap-

proach to social insurance laid out in section 3. Moreover, the empirical material presented in this section helps to identify those types of social insurance benefits that seem particularly well suited for finance via individual saving accounts. The present section therefore provides an important background for the specific reform proposal for Denmark presented in section 4.

2.1. Estimating lifetime incomes

To evaluate how the welfare state affects the distribution of lifetime incomes, one must construct measures of lifetime incomes and their distribution across the population in the absence and presence of taxes and transfers. For many years, the secretariat of the Danish Economic Council (DEC) has constructed such measures to highlight trends in the Danish income distribution.² The empirical analysis in this paper employs the most recent estimates of lifetime incomes in Denmark, presented in the council's report from the spring of 2005 (Danish Economic Council, 2005). This subsection explains the methodology used by the DEC; further details and documentation may be found in Hansen (2005).

The estimates are based on a comprehensive micro panel data set covering the period 1994-2002 and comprising a representative sample of 10 percent of the Danish population above the age of 18. The data include a wide range of socioeconomic variables plus information on annual factor incomes, public transfers received and taxes paid. The recorded factor incomes include an imputed return to owner-occupied housing, and incomes are measured in 2002 income levels (i.e., incomes for other years are adjusted for average nominal income growth), using an equivalence scale adjusting for economies of scale in the consumption of multiperson households.

The data cover a time span of nine years in the lives of the various cohorts aged 18 and above in 1994. Lifetime incomes are estimated by matching individuals from different cohorts with otherwise similar observable characteristics. This procedure implicitly abstracts from cohort effects other than those stemming from the observable socioeconomic

²Established by the Danish parliament in 1962, the Economic Council is an independent think tank advising the Danish government and parliament on issues of economic policy. The council is headed by three academic economists who prepare two reports on the state of the Danish economy every year, assisted by the professional economists in the council secretariat.

characteristics included in the data set. By combining the observations for different cohorts, one obtains synthetic life cycles covering the age interval from the age of 18 until the age of death of the oldest individual included in a given synthetic life cycle. The constructed life cycles thus exhibit different lengths, and the tendency for higher-educated people to live longer (and hence to benefit more from public pensions) is reflected in the data.³

The starting point for the construction of synthetic life cycles is the cohort aged 42 years in 1994 and thus 50 years in 2002. A person in this group with certain characteristics (Person 1) is matched with a person with similar characteristics who was 50 years old in 1994 (Person 2) in order to add observations of annual incomes in the age interval between 51 and 58. Similarly, Person 2 is matched with a person with similar characteristics who was 58 years old in 1994 (Person 3) to add another eight-year age interval to the constructed life cycle, and so on. Further, since Person 1 was 42 years old in 1994, he/she is also matched with a similar person who was 42 years old in 2002 (Person 4) in order to add observations for the age interval 34-41 years to the constructed life cycle, and Person 4 is in turn matched with a person who was 34 years of age in 2002, etc. This procedure means that a synthetic life cycle ending at the age of, say, 82 is constructed on the basis of data for eight different individuals, with the youngest one being 18 years of age in 1994 and the oldest one being 74 years of age in that year.

The procedure described above started from the cohort that was 50 years of age in 2002. A similar procedure is repeated eight times, each time starting with a cohort that was one year younger in 2002. The last set of synthetic life cycles is thus constructed by starting with those individuals who were 43 years old in 2002. Since each of these eight cohorts in the sample population includes more than 7,000 individuals, one ends up with more than 58,000 synthetic Danish life cycles. Centering the construction of life cycles around the cohorts aged 43-50 years in 2002 means that the resulting lifetime incomes reflect the current level of education of middle-aged Danes rather than the higher (lower) education level of younger (older) cohorts.

The purpose of the matching procedure is to ensure that the individuals who are linked together in the same life cycle are as similar as possible in terms of the socioe-

³The average age of death in the constructed life cycles is 75 years.

conomic characteristics determining lifetime income. Ideally one would like to match individuals who are fully identical with respect to gender, education, family status, sector of employment etc., and who have identical incomes at the same age level. However, although the sample population is large, such a matching procedure would imply a loss of a large number of observations due to missing matches, since in most cases it would be impossible to find individuals who are completely identical in terms of all observed characteristics, *including* the income they earn at a given stage in their life cycle. The matching of individuals is therefore carried out in two steps. The first step may be explained by going back to the cohort N_{02}^{50} of individuals who were 50 years old in 2002. Each of these persons needs to be matched with a similar person from the cohort N_{94}^{50} of people who were 50 years of age in 1994. For this purpose, all individuals within each of these two cohorts are divided into 60 different groups, categorized according to gender, three different levels of education, and ten deciles of annual disposable income. This initial categorization ensures a significant degree of similarity between individuals who are matched, since nobody from the cohort N_{02}^{50} can be matched with a person from the cohort N_{94}^{50} who belongs to another group.

In the second step, an individual from cohort N_{02}^{50} belonging to a given category X (Person 1) is matched with an individual from cohort N_{94}^{50} who also belongs to category X and who has an *expected* annual disposable income as close as possible to the income of Person 1 (recall that all incomes are measured in 2002 income levels and are thus directly comparable). The expected disposable income for a 50-year old in 2002 (1994) is estimated by running an OLS regression using data on all individuals who were in the age interval 50-54 years in 2002 (1994), incorporating 53 different socioeconomic characteristics as explanatory variables, including family composition, detailed level of education, employment status, ethnic background etc. In a similar way, the expected disposable income of, say, a 37-year old individual is estimated by running regressions on data of all individuals in the age interval 35-39 years in 1994 and 2002, respectively. Matching individuals on the basis of *expected* rather than *actual* incomes eliminates the effects of random fluctuations in individual incomes and allows the matching to exploit information on all the observable characteristics that tend to make the incomes of any two

individuals converge.⁴ At the same time, the categorization into 60 groups undertaken in the first step ensures that individuals who are matched always belong to the same decile in the distribution of *actual* incomes.

Note that while individuals are matched on the basis of their expected disposable income, the lifetime income in each synthetic life cycle is calculated from the actual observed annual incomes of the individuals included in the constructed life course. Further, although the matching is based on expected *disposable* income, the data set also allows one to track the evolution of actual *factor* income throughout each constructed life cycle. For simplicity, it is assumed that the relevant discount rate equals the average growth rate of real income (in recent years the average interest rate on government bonds has in fact been quite close to the rate of wage growth in Denmark). One may then simply add up the annual incomes earned in each constructed life cycle to obtain an estimate of lifetime incomes.

The final challenge is to estimate the distributional impact of the current tax-transfer system. The taxes and transfers recorded in each synthetic life cycle are influenced by policy rules dating back as far as 1994. To ensure that the taxes and transfers assigned to each life course reflect current rather than historical policy rules, we replace the recorded actual tax and transfer payments by the estimated tax-transfer payments that would have materialized in case the most recent policy rules would have prevailed throughout each individual life cycle. Specifically, the tax payments and transfers assigned to each synthetic life cycle are based on average observed payments for the years 2000-2002. For a person aged 50 years in 2002, this average payment is imputed to each of the years in the age interval 46-53 years in the synthetic life cycle in which he is included. To the age interval 54-61 years, one imputes the average annual amount of taxes and transfers recorded for 2000-2002 for the person in that same synthetic life cycle who was 58 years old in 2002; to the interval 38-45 years, one assigns the average 2000-2002 taxes and transfers for a person who was 42 years old in 2002, and so on. In this way, one obtains estimates of taxes and transfers over the entire life course, assuming that the tax-transfer

⁴This matching methodology is similar to the method of propensity-score matching, which has gained popularity in recent years as a means of matching treatment groups with appropriate control groups when evaluating the effects of various public-policy programs (see, e.g., Caliendo and Kopeinig, 2005).

rules in the period 2000-2002 prevail over the full life cycles of all individuals.⁵

Table 1 compares the estimated distribution of lifetime incomes with the observed distribution of annual incomes in 2002. Average factor incomes have been normalized to 100, so that all numbers in the table are measured in percent of average factor income. To facilitate comparison of the total incomes earned over life cycles of different lengths, and to allow comparison with the cross-section data on annual incomes reported in the upper part of the table, we divide all lifetime incomes by the number of years in the constructed life cycle in order to obtain a measure of the average income earned per year. Gross income is defined as the sum of factor income and the pre-tax public transfers received. Disposable income equals gross income minus direct taxes paid.

As might be expected, the distribution of annual income is much more unequal than the distribution of lifetime income. Measured by the reduction in the Gini coefficient, transfers have a greater equalizing impact than direct taxes, whether distribution is measured on an annual or on a lifetime basis. The difference between the Gini coefficients for the distribution of factor income and disposable income is sometimes taken as a measure of the total redistributive impact of the tax-transfer system. An obvious problem with this approach is that the distribution of factor income may be significantly affected by the tax-transfer system. With this important proviso in mind, we see from Table 1 that transfers and direct taxes cut the Gini coefficient for the distribution of lifetime income roughly in half. Interestingly, in relative terms, this reduction in inequality is slightly larger than the policy-induced 45 percent reduction in the Gini coefficient for the distribution of annual income. In absolute terms, however, the tax-transfer system has a larger negative impact on the Gini coefficient if income is measured on an annual basis rather than a lifetime basis.

(Table 1 about here)

⁵This procedure is necessitated by the fact that the period 2000-2002 only includes observations of tax/transfer payments during three years of each of the eight-year intervals making up a synthetic life cycle. Undoubtedly, our procedure implies an overstatement of the degree of persistence in individual tax payments and transfer receipts from one year to the next in the life cycle. However, over the course of an entire life cycle - which is the perspective adopted here - our procedure is unlikely to imply a systematic bias in the amount of taxes and transfers assigned to the various lifetime income deciles.

2.2. The redistributive impact of social-insurance transfers

Tables 2 through 4 describe the redistributive effects of the most important social-insurance transfers in the Danish welfare state. Based on cross-section data for 2002, Table 2 shows the fraction of annual disposable income accounted for by the various transfers in each of the deciles in the distribution of *annual* incomes. Table 3, in contrast, displays the fraction of disposable *lifetime* income accounted for by these same transfers in the various deciles of the lifetime income distribution, using the estimates of lifetime incomes presented in section 2.1.⁶

(Tables 2 and 3 about here)

Utilizing the information contained in Tables 2 and 3, Table 4 reports the redistributive impact of the various social-insurance transfers, measured by the 'redistribution index'. The point of departure for calculating the redistribution index is a standard Lorenz-curve diagram, in which the population is ranked according to income deciles along the horizontal axis and where the vertical axis measures the share of total income earned by the poorest X percent of the population. In addition to the Lorenz curve, such a diagram may include a *concentration curve* measuring the fraction of total spending on some social-insurance benefit accruing to the poorest X percent of the population. If people in poorer deciles receive a larger share of total spending on the transfer considered, the concentration curve will lie above the 45-degree line; in the hypothetical case of an identical lump-sum transfer to all citizens, the concentration curve will coincide with the 45-degree line. The redistribution index is defined as the area between the concentration curve and the Lorenz curve, measured in proportion to the total area below the 45-degree line. The greater the value of this index, the more redistributive is the transfer in question. Table 4 normalizes the redistribution index by reporting the *excess* value of the

⁶According to tables 2 and 3, early retirement benefits are a larger percentage of lifetime income than annual income in 2002, whereas the opposite is the case for child benefits and parental leave benefits. The difference with respect to early retirement benefits stems from the fact that the cohort of individuals who were between 43 and 50 year old in 2002 (the cohort around which our synthetic life cycles was centred) was significantly larger than the cohort that was entitled to early retirement benefits in 2002. At the same time, the 43-50 year olds in 2002 had a relatively low fertility rate, so in our synthetic life cycles these people collect a relatively small amount of child benefits and parental leave benefits.

index above the value of the redistribution index for an identical lump-sum transfer to all individuals. The index numbers in Table 4 thus indicate how much more redistributive the various transfers are compared to a uniform lump-sum transfer.

(Table 4 about here)

The first column in Table 4 shows the extent to which the most important social-insurance benefits help to redistribute *annual* incomes. In an annual perspective, social assistance benefits and education benefits are the most redistributive transfers. Also housing benefits and supplementary retirement benefits (which are means-tested) generate substantial redistributive effects. In a *lifetime* perspective, most transfer programs exert a smaller effect on the income distribution. Moreover, the ranking of the various transfers according to their redistributive impact changes significantly, as shown by the second column in Table 4. Social assistance remains the most redistributive program, but its redistributive effect is significantly smaller in a life-cycle context. Transfers such as parental leave benefits and the basic retirement benefit (which is a flat benefit granted to all Danish residents above the age of 65) that yield a significant impact on the distribution of annual incomes exert (almost) the same effect on the distribution of lifetime incomes as an identical lump-sum transfer to all individuals. Most strikingly, whereas education benefits are highly redistributive in an annual context, they generate only small effects on lifetime income distribution. Disability benefits, in contrast, yield a stronger redistributive impact in a lifetime perspective than in an annual perspective.

2.3. Interpersonal versus intrapersonal redistribution

Table 4 suggests that a large fraction of the taxes levied to finance public transfers serves to finance redistribution over the same individual's life cycle. But exactly how large is this fraction? To answer this question, we adopt a methodology previously used by Hussénius and Selén (1994), Falkingham and Harding (1996), O'Donoghue (2001), and Pettersson and Pettersson (2003). In applying this method, we use the constructed synthetic life cycles described in section 2.1 as an input. Our calculations account for all public transfers paid to Danish households. These benefits include all of those mentioned in Table 4 plus a few others that only make up a minor share of total transfers.

The first step in the calculation is to identify the fraction of taxes serving to finance social-insurance transfers. The total social-insurance benefits received over the life cycles of the representative sample of Danish households included in our data set amount to 43.7 percent of the total direct and indirect taxes paid by these same individuals over their lifetimes, with the remaining fraction of taxes serving to finance other categories of public spending. In our main scenario we therefore assume that 43.7 percent of direct and indirect taxes are 'reserved' for the financing of social-insurance benefits. The direct taxes paid by each individual in each year of his/her life cycle are explicitly included in our data set, so we take 43.7 percent of these payments to represent the financing of social insurance. Individual payments of indirect taxes are not explicitly included in our data, however, so these payments have to be estimated. According to the Danish Ministry of Finance (2002), the indirect taxes paid by the average Danish taxpayer amount to 22.3 percent of disposable income, varying from 37.4 percent of annual disposable income in the bottom income decile to 16.4 percent of annual disposable income in the top decile. We use these estimates to impute indirect tax payments to all individuals in our sample, accounting for the income decile to which they belong in each year of their life cycle. The fraction of income devoted to indirect taxes thus varies as individuals move from one place in the annual income distribution to another over the life cycle.⁷ Having allocated indirect taxes across individuals in this way, we assume that 43.7 of these payments are 'reserved' for the financing of social-insurance benefit.

We now wish to separate the *interpersonal* redistribution achieved via the social-insurance system (i.e. redistribution of lifetime incomes from rich to poor) from the *intrapersonal* redistribution (i.e. redistribution over the individual's life cycle). We therefore define the net lifetime transfer received by individual i over the life cycle as

$$N_i = \sum_t (B_{it} - T_{it}), \quad (2.1)$$

where B_{it} denotes transfers received by individual i in year t , and T_{it} is the part of individual i 's direct and indirect tax payment in year t that is 'reserved' for the financing of

⁷This procedure assumes that saving rates depend only on income. Our data set does not allow us to estimate how saving rates depend on other household characteristics. Below we shall consider a scenario where social transfers are assumed to be financed only out of direct taxes so that we do not have to make any assumptions regarding indirect tax payments.

social insurance. The total interpersonal redistribution achieved in a life-cycle perspective (*INTER*) is found by aggregating the net lifetime transfers across those individuals for whom the net transfer received is positive:

$$INTER = \sum_i (N_i | N_i > 0) = - \sum_i (N_i | N_i < 0). \quad (2.2)$$

The intrapersonal redistribution can be split into two components. Within any year in which an individual receives a transfer, he/she will also pay some amount of tax (at least indirect tax), so that part of the transfer received is self-financed. We define the part of the transfer financed by the taxpayer himself *within the same year* as

$$SY_{it} = \min(B_{it}, T_{it}). \quad (2.3)$$

Another part of the transfers received over the life cycle is financed by the recipient himself via the taxes paid in the other years of his life. This self-financing of benefits via the 'reserved' taxes paid in some other year may be calculated as

$$OY_i = \sum_t (T_{it} - SY_{it}) \quad \text{if } N_i > 0, \quad (2.4)$$

$$OY_i = \sum_t (B_{it} - SY_{it}) \quad \text{if } N_i \leq 0. \quad (2.5)$$

The total intrapersonal redistribution (*INTRA*) is the sum of the self-financing by all individuals within the same year and across their life cycles:

$$INTRA = \sum_i \sum_t SY_{it} + \sum_i OY_i. \quad (2.6)$$

We have now decomposed total transfer payments into interpersonal and intrapersonal redistribution, since the definitions given above imply that

$$INTER + INTRA = \sum_i \sum_t T_{it} = \sum_i \sum_t B_{it}. \quad (2.7)$$

Table 5, taken from Hansen (2005, p. 76), contains estimates of interpersonal and intrapersonal redistribution in Denmark, based on the synthetic life cycles described in section 2.1. All numbers in the table are average payments accumulated over the life cycle, accounting for observed differences in life expectancies across income groups. To put the numbers in the subsequent rows in perspective, the first row shows the average

lifetime factor income in the various deciles. The second row from the bottom of Table 5 shows the average net transfer received by those who end up with positive net receipts from the government over their life cycles (those for whom $N_i > 0$). In the bottom row, we report the average net taxes paid by those who end up paying more to the government than they receive ($N_i \leq 0$). Not surprisingly, most of the net recipients are concentrated in the lower lifetime income groups, while most of the net tax payments are concentrated in the top income groups, reflecting our earlier finding that some amount of redistribution from rich to poor does after all take place.

The first ten columns in Table 5 contain average accumulated lifetime payments/receipts for each of the ten income deciles. The final column reports averages across the entire population. Adding the figures in the fourth and fifth rows in the final column and dividing the sum by the figure in the second row of that same column, we find that 74 percent of the taxes levied to finance social insurance represent intrapersonal redistribution over the taxpayer's own life cycle, leaving only 26 percent of tax revenues for interpersonal redistribution from high- to low lifetime incomes. Moreover, it is striking that even at the bottom of the income distribution some people are net taxpayers on a lifetime basis, while at the top of the distribution some people receive net benefits over their life course. Presumably, an optimal tax-transfer system seeking to redistribute lifetime incomes would limit such cases.

(Table 5 about here)

The estimates in Table 5 assume that indirect as well as direct taxes contribute proportionally to the financing of social-insurance benefits. Since indirect taxes are typically levied for revenue purposes rather than to equalize the distribution of income, one might alternatively assume that social-insurance benefits are financed only out of direct taxes. Adopting this alternative assumption, Hansen (2005) finds that about 29 percent of social-insurance transfers serve to redistribute lifetime incomes in Denmark, compared to 26 percent when the financing includes indirect as well as direct taxes. The conclusion thus remains that by far the greatest part of redistribution is intrapersonal rather than interpersonal.

Assuming that both indirect and direct taxes contributed *pari passu* to the financing of social insurance, Hussénus and Selén (1994) employed the method described above

to estimate that only about 21 percent of the taxes levied to finance social insurance in Sweden accomplishes interpersonal redistribution. Their study was based on a much smaller panel data set covering only two calendar years. This necessitated the matching of a large number of individuals to construct synthetic life cycles. Since the statistical matching tends to reduce the recorded variation in lifetime incomes, the Swedish study may have underestimated the amount of interpersonal redistribution via the tax-transfer system.⁸ Pettersson and Pettersson (2003) recently updated and refined the estimates by Hussénius and Selén (op.cit.), estimating lifetime incomes with the aid of a dynamic micro-simulation model and including the value of important public services such as education, health care and care for the elderly in a comprehensive measure of lifetime income. With this extended concept of income, Pettersson and Pettersson found that only 18 percent of the taxes levied to finance social-insurance transfers and social services in Sweden can be categorized as interpersonal redistribution. Falkingham and Harding (1996) found a degree of interpersonal redistribution between 48 and 62 percent in Australia and between 29 and 38 percent in Great Britain, depending on the extent to which indirect taxes are assumed to contribute to the financing of social transfers. For Ireland and Italy, O'Donoghue (2001) estimated a degree of interpersonal redistribution of 45 percent and 24 percent, respectively, using the same assumption on indirect tax finance as the one employed in the present study and in the Swedish studies mentioned above.

These empirical findings suggest that the contribution of taxes and transfers to interpersonal redistribution is smaller in the Scandinavian countries than elsewhere. At least two factors may help to explain this difference, the first one being purely technical: in Scandinavia practically all social-insurance benefits are included in taxable income, whereas in other countries benefits are often tax-free (and correspondingly lower). The taxable status of benefits in Scandinavia automatically means that a relatively high fraction of social transfers will be financed by taxes paid by the benefit recipients themselves

⁸When individuals from two different cohorts need to be matched, some individuals from the smaller cohort are allowed to enter into more than one match in order to avoid the loss of too many observations. The fact that the same individuals are included in more than one synthetic life cycle tends to reduce the heterogeneity (and hence the dispersion of income) across individuals in the constructed life cycles. Compared to the study by Hussénius and Selén, this problem is smaller in the Danish study presented here where much fewer matches were required to generate a synthetic life cycle.

and will hence be categorized as intrapersonal redistribution. A second reason why the Scandinavian countries stand out may be that they tend to offer more universal benefits essentially covering the entire population, whereas social-insurance programs in the Anglo-Saxon countries tend to be more targeted towards the poor and to rely more on means-testing (see Lindbeck (2006)). Broader coverage of social transfer programs means that more taxpayers end up receiving substantial public transfers over their life course.

Still, the evidence indicates that even outside the Nordic area the benefits provided by the modern welfare state are to a large extent smoothing consumption across individuals' life courses rather than promoting equality in lifetime incomes or providing insurance against adverse shocks to lifetime incomes.

3. The theoretical rationale for social insurance based on individual savings accounts

As an alternative or supplement to tax-financed transfers, the consumption-smoothing function of the welfare state can be accomplished also through saving schemes that link taxes and benefits on an individual level. In such a scheme workers contribute a fraction of their earnings to an individual saving account that is debited when the owner draws social-insurance benefits. At the time of retirement, any surplus on the account is used to supplement retirement benefits. By linking benefits to contributions in an actuarially fair way, the saving accounts reduce the tax wedge on labour income. Social security contributions essentially become benefit taxes.⁹

With well-functioning capital markets, rational forward-looking behavior and no redistributional concerns, the government can rely on *voluntary* saving to accomplish consumption smoothing over the life course. *Compulsory* saving accounts can help address imperfect capital markets giving rise to liquidity constraints, just as they may help to address policy concerns regarding lifetime redistribution, lack of self control, and moral hazard in insurance. We explore these possible functions of compulsory saving accounts in turn.

⁹To the extent that existing social-security contributions finance wage-linked benefits, they are in fact already, at least in part, benefit taxes.

3.1. Liquidity insurance

Compulsory individual accounts protect people by allowing individuals to make withdrawals from the accounts even if the account balance is negative. In this way, the government in effect provides liquidity insurance and alleviates capital-market imperfections. The future compulsory contributions into the saving scheme act as collateral so that the government can provide credit. The relief of liquidity constraints is especially important for the lower middle class workers, who often face borrowing constraints that prevent them from smoothing consumption over time in the face of various shocks. By allowing workers to in fact access the capital market and thus decouple individual annual consumption levels from individual annual incomes, compulsory saving accounts make lifetime income rather than annual income a more important indicator of welfare.

In the context of a model with involuntary unemployment, Section A.2 in the appendix shows that compulsory saving accounts can offer more efficient liquidity insurance than tax-financed unemployment benefits do. Intuitively, by linking contributions more closely to actual benefits received, compulsory saving accounts contain the adverse incentive effects of providing liquidity insurance. Moreover, in contrast to a simple cut in unemployment benefits, the introduction of mandatory unemployment accounts and the associated liquidity insurance allows the individual to bear the risk of unemployment in a period of life when consumption is higher. Whereas a cut in unemployment benefits would force a cut in consumption when it is already low (assuming that the unemployed are liquidity-constrained and that benefits are lower than wages), achieving the same improvement of the public budget through mandatory contributions to individual unemployment accounts would allow workers to concentrate the cut in consumption in periods when they are employed and feature lower marginal utility of consumption.

3.2. Lifetime redistribution

In the presence of individual accounts, the government can protect the lifetime poor by bailing out individuals who end up with a negative account balance at the end of their working lives. In this way, the government redistributes to the lifetime poor and provides insurance against catastrophic shocks that substantially harm lifetime incomes. Redistribution is thus targeted more closely at the lifetime poor who are suffering a

combination of low wage incomes and frequent adverse shocks during their lives. The individual saving accounts essentially help to keep track of which individuals fare poorly in life. By thus in effect collecting information on who is lifetime poor, the individual accounts improve the equity-efficiency trade-off if the government uses this information to offer a lifetime income guarantee (in the sense of bailing out individuals who end up with a negative account balance). The government thus focuses its scarce resources on redistribution from the lifetime rich to the lifetime poor – rather than making politically expedient transfers among various important groups of voters with comparable long-run living standards. Indeed, by cutting out the transfers that merely redistribute resources over the life course and focussing the transfers on interpersonal redistribution to the lifetime poor, the government can reduce distortionary tax wedges on labor supply.

In the context of a model with three types of households, Section A.2 of the appendix shows how individual saving accounts can improve the equity-efficiency trade-off by using information on lifetime incomes. The additional information on lifetime incomes essentially allows the government to implement an optimal non-linear lifetime income tax. In particular, individual accounts establish an efficiency-enhancing actuarial link between contributions and benefits for high-income and middle-income workers – who currently pay distortionary taxes partly to finance distortionary social benefits to themselves – without reducing net transfers paid to the low-income workers who remain protected by the lifetime income guarantee. In other words, the saving accounts effectively enable the government to implement selective cuts in tax-financed benefits for high-income and middle-income groups without having to reduce these benefits at the bottom of the income ladder. By at the same time providing liquidity insurance, the government increases the importance of lifetime income rather than annual income as an indicator for overall welfare.

Individual accounts do not improve labour-market incentives for the lifetime poor. Indeed, the individual account system can be viewed as a way to implement low marginal tax rates at the top of the lifetime income distribution (see the model in section A.2 of the appendix). At the bottom of the lifetime income distribution, however, high marginal tax rates remain the inescapable price of redistribution. The government can rely on financial incentives to stimulate the middle class, which accounts for a large share of effective

labour input in the economy. However, the government must use other instruments to activate the lifetime poor, whose employment is important for maintaining social cohesion in a society. Among other things, the government can focus its active labour-market policies and its administrative resources on this group. In particular, the government may collect additional information by closely monitoring job search and health conditions. The government provides benefits on the condition that an able individual gives up leisure time to improve skills or (look for) work. In this connection, workfare may play a useful role; the mere threat of being put on workfare is likely to boost work incentives.

Individual accounts are particularly attractive if the distribution of life-course incomes is not very skewed compared to the distribution of annual incomes. In that case, annual income is typically not a good indicator for lifetime income; information on lifetime income can thus make lifetime redistribution more efficient. In the modern life cycle with many working women and long periods of full-time or part-time education, a substantial number of workers move between periods of full-time work to periods of voluntary (sometimes part-time) absence from the labor market to educate themselves, start of business, or care for children and/or frail relatives (see e.g. Bovenberg (2005)). This makes annual income a poor indicator of lifetime income.

With little lifetime inequality, redistribution of lifetime incomes does not have to be costly. Indeed, in that case, the government does not need to bail out many households with negative account balances. Intuitively, over their life cycles, a large middle class is able to finance its own benefits. This points to the importance of providing individuals with equal opportunities in the beginning of their working lives in the form of good start qualifications provided by basic education. The less polarized a society is in terms of human capital, the less the fiscal system has to redistribute resources from high lifetime-income earners to low lifetime-income earners and the more the government can limit itself to helping individuals smooth their consumption over the life cycle.

3.3. Mandatory saving and myopia

The individual accounts do not escape the trade-off between equity and efficiency to the extent that lifetime incomes are distributed unequally. Liquidity insurance also implies some costs. Lifetime redistribution as well as liquidity and lifetime income insurance

give rise to moral hazard; agents have an incentive to minimize their contributions and maximize their withdrawals. The government must therefore regulate withdrawals so that they can be made only for pre-specified purposes. Savings must also be mandatory – at least until a specific upper limit is reached.

In addition to moral hazard, lack of self control and myopia are other reasons for making saving mandatory. Compulsory saving accounts in effect extend mandatory saving aimed at retirement to precautionary saving aimed at social insurance for individuals of working age. By being paternalistic, the government helps individuals who lack self control to implement better consumption smoothing. However, if individuals lack the willpower or cognitive abilities to smooth consumption over their lifetimes, annual disposable income becomes relatively more important than lifetime income as a welfare indicator. Accordingly, the government should base its redistributive policies not only on lifetime incomes (and the associated balances in individual accounts), but also on disposable incomes at each point in time. Intuitively, the government cannot rely on individuals to allocate their lifetime incomes optimally over their life course, and must therefore be concerned also about the distribution of annual incomes.

A disadvantage of mandatory saving is that the government may force some people to save too much. This can be an important drawback if preferences are heterogeneous and people cannot undo mandatory saving by borrowing. In that case, tax incentives, which respect free choice, can complement limited mandatory saving as an instrument to stimulate individuals to save. Tax incentives, however, typically imply a large deadweight loss as individuals who would have saved even in the absence of tax incentives take advantage of the tax privileges by simply restructuring their portfolio. To prevent this, the government can target tax subsidies at agents with low financial and human wealth by limiting tax incentives to low levels of saving.

3.4. Moral hazard and optimal lifetime insurance

Another reason why saving schemes may enhance efficiency is moral hazard in insuring human-capital shocks over the life cycle. In particular, agents may be able to affect the probability that the insured contingency occurs. To illustrate, unemployment compensation can harm incentives to find work and remain employable. Another form of moral

hazard is benefit cheating, which can occur if the insured conditions are difficult to verify. Individuals may, for example, pretend to be sick or disabled in order to claim sickness or disability benefits. Moral hazard is a problem even for actuarially neutral insurances that charge a premium that is directly related to the expected individual benefit from the insurance.¹⁰

Various developments increase the dangers of moral hazard and hence make human-capital risks less insurable. As the economy shifts from blue-collar work in industrial sectors to white-collar work in service sectors and knowledge-intensive activities, mental causes of sickness and disability become more prominent. These types of sickness and disability can be less easily verified than physical disabilities. Moreover, an increasing number of workers now moves between periods of full-time work to periods of voluntary absence from the labor market to enjoy leisure, educate themselves, set up a business, or care for children or frail relatives. In such a transitional labor market, it becomes more difficult to separate voluntary periods of inactivity from involuntary unemployment. At the same time, individuals can increasingly affect the probability that they become unemployed by investing in their own employability. In other words, the dividing line between the contingencies that people are responsible for and these that they are not becomes less clear. These changes in the nature of social risks make it more costly to insure human capital in terms of harming the incentives to accumulate and maintain that capital. At the same time, a more dynamic world economy and a decline of the extended family as an insurance device have increased the demand for such insurance as people experience more substantial economic insecurity.

Moral hazard gives rise to a fundamental conflict between facilitating insurance and providing incentives to reduce the probability that the insured risk occurs. In particular, reducing the extent of insurance through the introduction of deductibles can combat moral hazard. Deductibles help internalize the social costs of benefit payments, thereby discouraging individuals from making excessive claims on the welfare state. At the same time, however, these deductibles impose costs on a risk-averse individual by reducing insurance through risk pooling. Another way to combat moral hazard is to monitor

¹⁰Whereas the premia of such insurances do not distort the labor market, the benefits harm labor-market incentives if people can affect the probability that they are eligible for these benefits by changing their labor-market behavior.

agents and to regulate their behavior, but this may well be costly in terms of intrusion in private lives.

Individual saving accounts can improve the trade-off between insurance and incentives by facilitating self-insurance over the life course. In particular, these accounts increase the scope for deductibles without compromising minimum consumption standards of individuals who are hit by temporary adverse shocks. They do so by allowing individuals who suffer from liquidity constraints when they are hit by an adverse shock to shift the payment of deductibles to the periods in which they can more easily afford these costs. Individuals can thus self-insure themselves over their life course, and do not have to rely on insurance that gives rise to moral hazard. Risks can be self-insured on a lifetime basis, and thus do not have to be insured on a day-to-day basis. Indeed, risks that may seem large on an annual basis may in fact be quite small when considered over an entire lifetime. To illustrate, two unemployment spells of half a year reduce lifetime incomes of an individual with a full-time working career of thirty years by only about 3 %.

The potential of individual accounts in improving the trade-off between insurance and incentives depends crucially on the extent to which individuals face correlated shocks during their lifetimes. The potential welfare gains of individual saving accounts are large if various income shocks are uncorrelated across time and among each other. In that case, annual incomes are poor indicators of lifetime incomes, and income shocks are in fact only small in the context of an entire lifetime. There is thus ample scope for self-insurance by pooling risks facing a single individual. If shocks are strongly positively correlated, in contrast, risks do not become much smaller in a lifetime context (compared to an annual context). In particular, some individuals are always unlucky and therefore remain poor, while others seem to continuously strike it rich. Risks then remain catastrophic, even when viewed over the entire life course. Self-insurance is then costly, and pooling risks across individuals (rather than just intertemporally for each single individual) through insurance creates substantial value. Also the scarring effect of unemployment on human capital makes insurance more valuable. More generally, labor-market risks tend to be correlated in the presence of dual labour markets in which insiders enjoy high incomes throughout their lives while outsiders must make do with insecure jobs and tend to suffer from frequent and long-lasting unemployment. Hence, long unemployment durations in

slow-moving labor markets make individual accounts less attractive as an instrument to provide lifetime income insurance.

For each type of human capital risk, another combination between insurance and self-insurance through saving is optimal, depending on the magnitude of the risk in terms of the potential drop in lifetime income and the potential danger of moral hazard because of endogeneity and non-verifiability of the insured risk. Self-insurance should be relatively important for non-catastrophic risks that people can affect through non-verifiable actions, such as short-term unemployment and the first sickness days. Stiglitz and Yun (2002) explore the optimal mix of self-insurance through saving accounts and tax-financed insurance. They show that self-insurance should play a more prominent role if risk aversion is low, moral hazard is important, various risks are uncorrelated across time and among each other, and these risks are only small in a lifetime perspective. They also demonstrate that the optimal extent of self-insurance depends on the history of an individual. Self-insurance should optimally be the most important for those individuals who have not experienced adverse shocks early in life so that they are not likely to end up being lifetime poor. Also here, the conclusion is thus that saving schemes can play a more important role in providing incentives for the middle- and higher incomes than for the lifetime poor. The optimal mix between saving and insurance may also vary between workers in different sectors in the economy. This provides an argument for a role for social partners on a sectoral level in determining the optimal mixes between saving and insurance. Indeed, individual saving schemes may be incorporated in collective sectoral agreements. These agreements may provide for mandatory contributions into both specific employee insurances with deductibles and individual saving schemes from which individuals can draw to pay deductibles.

Bovenberg and Sørensen (2006) investigate the optimal structure of lifetime income taxation and social insurance aimed at both lifetime income redistribution and disability insurance, which can be interpreted as insurance against all kinds of idiosyncratic shocks to human capital. They show that even in the absence of moral hazard full insurance against these idiosyncratic shocks is not optimal. The reason is that imperfect insurance encourages workers to self-insure themselves by raising their labour supply, thereby alleviating the distortionary impact of redistributive labour taxation. Hence, a tension exists

between lifetime redistribution and insurance against human-capital shocks. By harming labour supply, social insurance imposes a negative externality on the redistributive branch of the government. The greater the weight attached to redistribution towards the lifetime poor with low skills, the less the government can allow the higher skilled to insure themselves fully against disability risk. This result suggests that the government cannot leave individuals completely free to use their social security contributions to buy actuarially fair insurances. Against this risk of overinsurance, however, stands the risk of underinsurance due to selection giving rise to excessive transaction costs. In any case, a hybrid system of insurance and self-insurance would typically be optimal, depending on moral hazard, selection, risk aversion and redistributive preferences and the associated distortionary taxes.

3.5. Extensions

In outlining the case for individual saving accounts, we have focussed on insuring shocks affecting the earning capabilities of individuals. Similar arguments for and against saving schemes apply to the financing of health care. In particular, saving schemes can be part of a three-pillar model in health-care financing involving a hybrid system of saving, insurance and redistribution. In particular, this model involves, first, government assistance for those who cannot afford a minimum level of medical care (i.e. redistribution); second, medical insurance for catastrophic events supplemented by limited insurance for other events (i.e. insurance based on risk pooling); third, compulsory individual medical saving for financing deductibles and coinsurances (i.e. consumption smoothing and self insurance). The optimal mix between these three pillars depends on the particular type of health-care cost considered. As explained above for social insurance, saving schemes are most attractive for costs that are distributed rather uniformly across individuals (seen over the life cycle as a whole).

The principle of individual saving accounts can be applied to finance user fees for not only medical care but also other services, such as higher education and child care. If individuals pay these costs from their individual saving accounts and can thus smooth these costs over their entire lifetime, the government can rely more on consumer sovereignty for selecting the level, quality and nature of the service, and thus does not have to impose

strict regulations (e.g. by rationing individuals).

4. Individual accounts in a Western European welfare state: how would they work?

Having laid out the theoretical rationale for social insurance based on individual accounts combined with a lifetime income guarantee, we will now discuss in more detail how individual accounts (IAs) might be designed in practice, and how they are likely to affect income distribution, economic efficiency and the public finances. For the sake of concreteness, we consider the proposal of the Danish Economic Council (2005, ch. VI) for a system of individual accounts in Denmark. Section 4.1 describes the proposal and discusses some policy choices involved in its design. Section 4.2 estimates the effects of the proposed IA system on the distribution of lifetime incomes, and Section 4.3 explores its likely effects on the labour market, the public budget and economic efficiency. Section 4.4 compares our proposal to other recent proposals for an IA system.

4.1. A reform proposal for Denmark

As documented in section 2.3, there is a large element of intrapersonal redistribution in the current Danish welfare state arrangements. Against this background, the Danish Economic Council (2005, ch. VI) proposed that part of the existing social-insurance benefits should be financed via mandatory individual savings accounts. The purpose of the reform proposal (henceforth the DEC proposal) was to reduce the distortionary impact of the tax-transfer system while preserving the liquidity insurance and lifetime income insurance offered by the present system. To understand the context for the DEC proposal, note that in Denmark the bulk of social-insurance benefits is financed out of general tax revenues, and most benefits are paid out in flat rates that are unrelated to previous wage incomes. The link between benefits and labour supply is thus very weak in the current Danish system of social insurance. This implies that saving accounts can not only combat moral hazard in social insurance but also reduce the implicit tax wedge on labor supply due to a weak link between contributions and expected benefits.

According to the DEC proposal, the IA system would have the following features.

Each citizen in the age group from 18 years until the official retirement age of 65 years would be required to deposit a certain percentage of his/her labour income in an individual account every year. For employees, this social security contribution would be calculated as a percentage of gross wage income. For the self-employed, it would be calculated on the basis of the imputed labour income, which is computed every year as an integral part of the Danish tax code. Whenever a person receives certain social-insurance benefits according to the current eligibility rules (which are assumed to be unchanged), his/her individual account would be debited by the corresponding amount, and the balance on the IA would be carried forward with the interest rate on short-term government bonds. If the IA balance is positive at the time of retirement, the account holder may choose to have it converted into an annuity that is added to the ordinary public retirement benefit, or he/she may choose to have the balance paid out as a lump sum.¹¹ If the IA balance is negative when the individual reaches the official retirement age, the account is simply set to zero. Accordingly, the owner of the account still receives the ordinary flat retirement benefit as a consequence of the lifetime income insurance built into the system. For married couples, any benefit paid to one of the spouses is debited by half the amount on the IA of each spouse, and for unmarried parents any child-related benefits are likewise debited by half the amount on the IA of each parent. These rules are intended to ensure a reasonably equal distribution of IA balances between men and women.

When selecting the transfer programs to be included in the IA system, the DEC focused on those programs that involve the lowest degree of interpersonal redistribution in order to minimize the potential negative impacts on lifetime income distribution. Specifically, the DEC proposed inclusion of the following transfers in the IA system:

1. Early retirement benefits
2. Education benefits
3. Short-term unemployment benefits

¹¹When IA balances are converted into annuities, policymakers may choose to differentiate the conversion factors across different groups to reflect differences in expected lifetimes. As a technical matter, the calculations presented in sections 4.2 and 4.3 assume that all positive IA balances are paid out as lump sums at the time of retirement.

- (for unemployment spells up to a length of three months)
4. Sickness benefits (up to a limited number of sickness days)
 5. Child benefits
 6. Parental leave benefits

Table 4 shows that early retirement benefits, education benefits, parental leave benefits and child benefits imply a low degree of lifetime income redistribution, serving mainly to redistribute resources across the taxpayer's own life cycle. The same goes for ordinary retirement benefits, but this transfer program was not included in the proposed IA system in view of the ongoing discussion in Denmark about the need for a separate reform of the pension system in the light of population aging. Programs such as social assistance, disability benefits and housing benefits were excluded from the IA system since they involve a high degree of interpersonal redistribution, as indicated in Table 4. Furthermore, according to estimates by Hansen (2005, p. 86), the degree of lifetime income redistribution implied by benefits paid to workers suffering long unemployment spells (exceeding three months) is almost twice as large as the interpersonal redistribution generated by short-term unemployment benefits (for spells shorter than three months). For this reason the DEC proposal includes only short-term unemployment benefits in the IA system. Similarly, there is a presumption that benefits paid during long sickness spells are more redistributive than those paid during short spells. Moreover, short-term sickness spells tend to involve a greater moral hazard problem of verifiability. The DEC therefore proposed that only benefits paid during a limited number of sickness days should be included in the IA system. However, data limitations compelled us to include all sickness benefits in the calculations presented below.

The total spending on the transfer programs included in the proposed IA system amounts to 7.9 percent of the base for the proportional Danish payroll tax ('arbejds-markedsbidrag'). This tax is levied on gross wage income and on the imputed labour income of the self-employed (with no ceiling for any of these groups); for wage earners, the tax is collected at the employee level. The DEC therefore proposed that the payroll tax be cut by 7.9 percentage points and be replaced by a corresponding mandatory contribution to the taxpayer's individual account. This IA contribution would be deductible against the base for the personal income tax, as is the case for the current payroll tax. For

symmetry reasons, the positive balance on the IA would then be included in the personal income tax base when it is paid out. For individuals who end up with a negative IA balance, the deductibility of contributions is likewise offset by the fact that the benefits included in the IA system are all subject to income tax in the year they are paid out.

The account system could be administered directly by the government, or it could be administered by private sector financial institutions. In Denmark, all taxpayers below 65 years of age already contribute to a mandatory supplementary pension scheme (denoted ATP), so it would seem natural to build on this existing administrative infrastructure. The IA system would be phased in gradually. All citizens in the age group between 18 years and the official retirement age (currently 65 years) could start participating immediately. In the first few years, people reaching the retirement age would have accumulated relatively small balances. As time elapses and new retirees would have participated in the IA system for a longer time, balances would grow, but at the same time the labour-supply effects discussed below would exert a growing positive influence on net public revenue. The system is thus automatically phased in gradually, with no need for special transition rules.

4.2. Effects on income distribution

Table 6 shows the estimated distributional impact of the proposed IA system, assuming a zero growth-adjusted real interest rate. The estimate was produced by Hansen (2005) using the data and the constructed synthetic life cycles described in section 2.1. Importantly, the table abstracts from any behavioral effect of the IA system. The numbers thus reflect only the mechanical impact effect. Although the very purpose of the IA system is to influence labour-market behavior, the distribution of positive IA balances in Table 6 should provide a good proxy for the effect of the reform on the distribution of individual welfare. The reason is that, by the Envelope Theorem, changes in employment caused by the IA system yield no first-order effects on individual welfare if individuals have optimized their behavior in the initial equilibrium and are not rationed on the labor market.

(Table 6 about here)

The second row in Table 6 shows the accumulated contributions to the IA relative to the accumulated withdrawals from the account for each of the deciles in the lifetime income distribution. Not surprisingly, this ratio is systematically rising with lifetime income. Moreover, the ratio of the average positive IA balance to lifetime income is also rising with the income level, as shown in the third row in Table 6. Furthermore, whereas only 7.2 percent of individuals in the lowest decile end up with a positive IA balance at the time of retirement (assuming unchanged behavior), almost 80 percent of people in the top decile will accumulate a positive balance, as indicated in the fourth row of the table.

This distributional pattern reflects the fact that the contributions to the IA are proportional to labour income whereas most of the benefits included in the IA system are paid out in flat rates. It also reflects the fact that people who are less active in the labour market and more dependent on the transfer system tend to end up in the lower lifetime income brackets. There is thus no doubt that the proposed IA system will make the lifetime income distribution more unequal. The distributional impact will be limited, however. Specifically, while the Gini coefficient for the distribution of disposable lifetime income is 0.127 under the current Danish tax-transfer system (see Table 1), it would only rise to 0.133 if the proposed IA system were introduced. Table 1 reveals that the Gini coefficient for the distribution of lifetime factor income is currently 0.253. While the redistribution of lifetime income implied by the current tax-transfer system amounts to $(0.253-0.127)/0.253 = 49.8$ percent, the redistribution under the DEC proposal would thus still amount to a substantial $(0.253-0.133)/0.253 = 47.4$ percent.¹² Moreover, as we shall argue in the next section, the proposed IA system would generate a Pareto improvement even under rather conservative assumptions regarding behavioral responses.

4.3. Effects on the public budget and economic efficiency

Although the rate of contribution to the IAs is chosen so that total contributions correspond to total spending on the relevant transfers, part of the contributions are channeled back to the contributors in the form of positive IA balances. In the absence of any be-

¹²These mechanical calculations are based on the heroic assumption that factor incomes are unaffected by the tax-transfer system.

havioral changes, the positive IA balances recorded in the third row of Table 6 would thus imply a corresponding deficit on the public budget.

However, the IA reform is bound to affect labour supply through several channels. First, the reform implies a cut in the effective *marginal* tax rate on labour income for all those who can look forward to a positive IA balance. The reason is that a marginal contribution to the IA, which replaces the payroll tax, is returned to the taxpayer himself at the time of retirement (with interest added). The lower marginal tax rate should stimulate labour supply at the *intensive* margin, inducing people to work more hours. Second, by cutting the payroll tax, the reform also reduces the *average* effective tax rate on labour income, thus increasing the income gain experienced by a person who moves from non-employment into employment. This will boost labour supply at the *extensive* margin, thus increasing the rate of employment. Third, in the transfer programs included in the IA reform, the effective benefit rates will be reduced to zero, since an increased take-up of benefits is matched by a corresponding reduction of future retirement benefits (in present-value terms) for individuals with positive IA balances. This drop in the effective replacement rate in the relevant transfer programs also stimulates labour supply at the extensive margin.

An ideal tool for the evaluation of these behavioral responses would be a disaggregated econometrically estimated computable general equilibrium model. However, since no such model is yet available for the Danish economy, we resort to some aggregate back-of-the-envelope calculations. Section A.1 of the appendix shows that, if the growth-adjusted real interest rate is zero and income effects on labour supply are negligible, the effect of the IA reform on the public budget is given by

$$\frac{dR}{ewh} = \overbrace{d\tau + c \left(\frac{1-e}{e} \right) d\alpha}^{\text{mechanical effect on revenue}} - \overbrace{\varepsilon \left(\frac{\tau}{1-\tau} \right) d\tau - \eta \left(\frac{t+c}{1-t-c} \right) dt + c\eta \left(\frac{t+c}{1-t-c} \right) d\alpha}_{\text{revenue change due to labour-supply response}}, \quad (4.1)$$

$$\varepsilon \equiv \frac{dh/h}{dw(1-\tau)/w(1-\tau)}, \quad \eta \equiv \frac{de/e}{dy/y},$$

where dR/ewh is the present value of the change in the budget balance relative to the labour income tax base for individuals with positive IA balances, τ is the total marginal effective tax rate on labour income (including indirect taxes), c is the average replacement rate in the transfer programs included in the IA system (measured as the after-tax

benefit rate relative to the average pre-tax income of a full-time worker), e is the rate of employment (the participation rate), t is the total average effective tax rate on labour income (including indirect taxes), h is the number of hours worked by the average worker, w is his pre-tax real wage rate, ε is the elasticity of hours worked with respect to the marginal after-tax wage rate, y is the gain in real disposable income when moving from non-employment into employment, η is the elasticity of labour force participation with respect to y , and α is the fraction of the total transfers received by the average worker that is debited to his individual account.¹³

The first two terms on the right-hand side of (4.1) represent the 'mechanical' effect on revenue in the absence of any behavioral changes, while the last three terms capture the improvement of the budget generated by the labour-supply responses to the reform. Note that since α is zero before the reform, and since the reform involves a cut in the effective marginal and average labour income tax rate, we have $d\alpha > 0$, $d\tau < 0$ and $dt < 0$. Hence, all of the last three terms in (4.1) contribute to an improvement of the public budget on the assumption that the elasticities ε and η are positive. The budgetary effects of the IA reform arise only from its impact on the group that ends up with positive IA balances. According to the bottom row in Table 6, a lower-bound estimate of the size of this group is that it will include about 46 percent of all taxpayers. However, Table 6 abstracts from behavioral changes. In practice, these changes enable some taxpayers who would otherwise have ended up with a small IA deficit to accumulate a surplus. The Danish Economic Council (2005, ch. VI) therefore estimates that about 60 percent of all taxpayers will end up with a positive IA balance. The remaining 40 percent of taxpayers with negative balances are treated exactly as under the current tax-transfer system. This group may therefore be neglected when evaluating the budgetary effects of the reform.

Since the cut in τ depends on the amount of benefits that are financed via the IAs, there is a systematic link between the magnitudes $d\tau$ and $d\alpha$ in (4.1). Specifically, section A.1 of the appendix shows that

$$d\alpha = - \left(\frac{ewh}{b(1-e)} \right) \left(d\tau + \frac{A}{ewh} \right), \quad (4.2)$$

where b is the average after-tax benefit rate in the transfer programs included in the

¹³With negligible income effects, the elasticities ε and η reflect compensated as well as uncompensated elasticities.

IA system, A is the average IA balance for those with a positive balance at retirement (assuming unchanged behavior), and ewh is the average lifetime labour income in the group of taxpayers with a positive IA balance (also assuming unchanged behavior). Given the parameter values reported in section A.1 of the appendix, plausible parameters in a Danish context would be

$$\frac{ewh}{b(1-e)} = 20.39, \quad \frac{A}{ewh} = 0.016. \quad (4.3)$$

In the DEC proposal, the payroll tax rate is cut by 7.9 percentage points, but since the payroll tax is deducted from the income tax base, the fall in the effective marginal direct tax rate on labour income is only about half this amount, given the estimate by the Danish Ministry of Finance (2004) that the marginal income tax rate for the average Danish taxpayer is about 50 percent. Further, using formula (A.8) in the appendix, we find that a fall in the effective marginal *direct* tax rate of about 4 percentage points translates into the following drop in the *total* effective marginal tax rate (which includes indirect taxes): $d\tau = -0.029$. Substituting this along with (4.3) into (4.2), we get $d\alpha = 0.263$. The interpretation of this number is that on average about 26 percent of total transfers are debited to the IAs, reflecting the fact that the IA system applies only to a subset of social-insurance programs (where 100 percent of benefits are debited to the IAs).

According to section A.1 of the appendix, we initially have (roughly)

$$t = 0.54, \quad \tau = 0.63, \quad c = 0.23 \quad (4.4)$$

for the average Danish worker. Moreover, formula (A.8) in the appendix implies that $dt = d\tau = -0.029$. To be able to apply formula (4.1), we now need only to calibrate the labour supply elasticities. Although estimates of the average (uncompensated) wage elasticity of hours worked for Denmark tend to centre around 0.1 (a little higher for females and a little lower for males), we select $\varepsilon = 0.05$ to be on the safe side. The participation elasticity η was recently estimated by Le Maire and Scheuer (2005) to be in the range 0.2-0.4 for Danish recipients of social assistance benefits. However, the authors argue that these estimates may have an upward bias, so to remain conservative we set $\eta = 0.1$. This relatively low estimate partly reflects the effectiveness of Danish active labor-market policies in encouraging transfer recipients to find work.

Armed with all these parameters, and noting that $c(1 - e)/e \equiv b(1 - e)/ewh = 1/8.5 = 0.118$, we now find from (4.1) that

$$\frac{dR}{ewh} = \overbrace{-0.016}^{\text{mechanical effect on revenue}} + \overbrace{0.002}^{\text{revenue effect of hours-of-work response}} + \overbrace{0.01}^{\text{revenue effect of participation response to lower taxes}} + \overbrace{0.02}^{\text{revenue effect of participation response to lower net benefits}} = 0.016. \quad (4.5)$$

Despite our rather conservative labour supply elasticities, the IA reform would improve the public budget by about one and a half percent of the labour-income tax base for the individuals whose incentives are positively affected. Indeed, we see that the labour-supply response to the cut in the marginal and average effective labour-income tax rates would in itself enable the government to recoup $(0.002 + 0.01)/0.016 = 75$ percent of the initial revenue loss from the reform.¹⁴ (4.5) shows that a substantial part of the positive impact on public revenue stems from the participation response to the cut in effective benefit rates implied by the IA system. On reflection, this is not surprising, given that the IAs effectively reduce the replacement rates in the relevant transfer programs from around 50 percent to zero, and that the effective tax wedge on the participation margin $\left(\frac{t+c}{1-t-c}\right) = 3.34$ is quite large.

If workers were allowed to use their accounts to buy actuarially fair insurance from private insurance companies (for example involving sickness or disability insurance), the actuarial link between contributions and expected benefits would reduce tax distortions compared to the present system of social insurance, but effective benefit rates would remain constant. Accordingly, the behavioral responses would be limited to those stemming from cuts in marginal and average tax rates.

The finding in (4.5) implies that the IA reform would be a genuine Pareto improvement. The reform raises the welfare of all agents who end up with positive IA balances (and thus enjoy higher retirement benefits), generates a bit of additional revenue for the government, and leaves the agents who end up with negative IA balances unaffected.

¹⁴The initial revenue loss from the reform consists of a loss of 2.9 percent of the tax base due to the tax cut, and a revenue gain of 1.3 percent of the tax base due to the cut in the effective benefit rates. The labour supply response to the lower marginal and average tax rates thus allows the government to recoup $0.012/0.029 \approx 41$ percent of the revenue loss from the tax cut. By way of comparison, using a computable general equilibrium model, the Danish Economic Council (2004, p. 94) estimated that about 56 percent of the initial revenue loss from a cut in the payroll tax rate would be recouped via increased economic activity. Our estimate is thus slightly more pessimistic than that of the DEC.

The Danish Economic Council (2005, ch. VI) reaches a similar conclusion. It likewise estimates that the reform would slightly improve the government budget, using a somewhat different method of calculation. Furthermore, the result that the introduction of IAs would imply a Pareto improvement is in line with the theoretical analysis in section A.2 of the appendix.

One may still wonder whether the result reported in (4.5) is too good to be true. In particular, agents respond to the incentives provided by individual accounts only if they are forward-looking. Empirical evidence suggests that some agents are in fact myopic. For example, in a panel study of consumption behavior in five industrial countries, Campbell and Mankiw (1991) estimate that only between 60 and 80 percent of consumers are forward-looking. If we multiply the assumed labour-supply elasticities in (4.1) by a factor of 0.6 as an ad hoc way of accounting for myopic behavior, we obtain $dR/ewh = 0.003$. Thus, even with a significant degree of myopia, the reform would still be fully self-financing.

It may be argued that, by neglecting income effects, our analysis tends to overstate the positive effect of the reform, since the positive IA balances will exert a negative income effect on labour supply. However, one should take into account that we have also abstracted from the fact that lower effective tax and benefit rates will tend to reduce structural unemployment and may stimulate human capital formation. Moreover, recent empirical studies suggest that income effects on labour supply are very small indeed. Overall then, it does not seem at all farfetched to claim that the introduction of IAs has the potential to generate a Pareto improvement.

The decomposition of the budgetary impact in (4.5) allows a quantification of the efficiency gains from the introduction of IAs. The revenue generated by the labour-supply response to the reform is roughly equal to the increase in labour supply times the tax and benefit wedge between the marginal productivity of labour and the marginal disutility of work. To a first-order approximation, this product reflects the efficiency gain from the reform if we abstract from involuntary unemployment.¹⁵ It is given by the sum of the last three figures on the right-hand side of (4.5), amounting to 3.2 percent of the

¹⁵In the presence of structural unemployment, an increase in employment increases welfare by more than the additional government revenues on account of a broader tax base and a narrower benefit base. Hence, the welfare gains as computed here provide a lower bound for the actual welfare improvements.

tax base. Another way of decomposing the welfare gain is to note that, by the Envelope Theorem, the effect of the reform on individual welfare is equal to minus the mechanical effect on revenue, i.e. 1.6 percent of the tax base. In addition, the reform improves the budget by 1.6 percent of the tax base, implying a total welfare gain of 3.2 percent of the tax base if the extra revenue is channelled back to the private sector in a lump-sum manner.

4.4. Alternative schemes

The saving accounts we have simulated in this paper are designed to ensure that nobody loses from the accounts ex post. Hence, we take the status quo as the starting point and explore whether we can establish a Pareto improvement, not only from an ex-ante but also from an ex-post perspective. Others have proposed other types of individual saving accounts. Just as we do, Fölster (1999) takes the currently paid taxes as the point of departure, reducing taxes by an amount corresponding to the mandatory contributions to the IAs. However, Fölster's scheme does not guarantee that nobody loses from the saving system ex post. Instead, he calibrates the minimum guaranteed pension so as to ensure that total pensions paid out under the account system (including the positive IA balances) equal the total amount of pensions paid out under the present system. This means that individuals who have a relatively high wage towards the end of their career (and who would therefore be entitled to a relatively high pension under the current system) but who nevertheless end up with a negative IA balance will tend to lose from the reform, since they are entitled to only a relatively low minimum pension. Another difference with the accounts we simulate is that Fölster (1999) finances the income guarantee through an explicit insurance premium payable by everybody. As a consequence, marginal tax rates remain positive for individuals who do not expect to benefit from the income guarantee.

Orzag and Snower (2002) focus on unemployment accounts. Workers would be required to put mandatory contributions in their unemployment account to finance withdrawals when unemployed. Individuals with a zero balance would be entitled to unemployment assistance. In addition, the contributions of workers earning low incomes could be subsidized. These subsidies and unemployment assistance would be financed by taxes on the contributions of other workers. In this way, marginal taxes remain positive for

all. Compared to the system we simulate, the tax rate on incomes is less non-linear so that also the lifetime poor may face lower marginal tax rates and improved incentives to contain moral hazard. At the same time, marginal tax rates remain positive also for those earning middle incomes and higher incomes. Moreover, some lifetime poor individuals as well as individuals who experience several bouts of unemployment during their life course may lose from the introduction of individual accounts.

Stiglitz and Yun (2002) explore social-insurance accounts that are integrated with the retirement system. Whereas the accounts raise expected utility ex ante, they do not guarantee that all agents are better off ex post. In particular, agents who have suffered frequent and lengthy spells of unemployment during their lives may end up with lower retirement benefits than with unemployment benefits that are entirely financed through taxes. Within this setting, Stiglitz and Yun (2002) explore which share of unemployment benefits should be optimally financed from saving accounts from the point of view of maximizing ex-ante utility of the workers. They argue that explicit transfers from high-skilled to low-skilled workers could mimic the transfers between these groups implicit in the current unemployment insurances without undoing the efficiency gains from individual accounts in combatting moral hazard.

Leijnse et al. (2004) propose a three pillar system. The first pillar resembles the transfer programs we do not include in the IA system. It remains completely tax financed. The third pillar includes voluntary saving schemes. The second pillar involves mandatory contributions to individual accounts that offer neither a lifetime income guarantee nor liquidity insurance. As a direct consequence, some individuals who make frequent withdrawals and exhaust their accounts would lose ex post. The second pillar, however, would be a mixture between insurance and saving. Hence, on a sectoral level, social partners can force workers to use part of their contributions to buy insurance. In that case, potential ex-post losses would be contained, as would be the potential welfare gains from reduced moral hazard. Indeed, the optimal mix between saving and insurance would depend on the scope for moral hazard. In particular, the share of saving and self-insurance would increase if individuals bear a larger personal responsibility for an event.

In summary, there are many possible ways of designing a system of individual accounts for social insurance. Comparing the costs and benefits of alternative designs is an

interesting topic for future research, but a systematic comparison goes beyond the scope of this paper.

5. Concluding remarks

Our analysis suggests that individual accounts can play a useful role in financing social benefits that have only little redistributive power in a life-cycle perspective and give rise to serious moral hazard. For such benefits, saving accounts can enhance labor-market incentives at a relatively low cost in terms of a more unequal distribution of lifetime incomes. This is especially so if saving accounts are accompanied by labor-market institutions that combat long-term employment and facilitate rapid turnover and by social policies that provide a life-time income guarantee and ensure an equal distribution of human capital at the beginning of life. As the changing nature of social risks makes social insurance more expensive in terms of distorted labor-market incentives, individual accounts with a lifetime income guarantee seem to be an attractive alternative to simple cuts in tax and benefits. Indeed, such accounts can continue to provide substantial income security at a time when a dynamic world economy confronts many people with substantial risks. In this way, they can help protect the social legitimacy of a competitive market system that stimulates innovation and growth but also gives rise to substantial risks associated with creative destruction.

Apart from the changing nature of social risks and the continued demand for income security, several factors have made individual accounts in social insurance more attractive. First of all, modern information and communication technologies enable governments to keep systematic records of the contribution and withdrawal histories of their citizens. Second, more efficient capital markets allow individuals to smooth their consumption over their life courses. By thus allowing individuals to decouple annual consumption from annual disposable incomes, better functioning capital markets make lifetime- rather than annual incomes better indicators of overall welfare. Moreover, financial innovation allows private financial institutions to administer the compulsory saving accounts. A further reason for the increased attractiveness of individual accounts is that they are fully portable between jobs. Hence, social insurance does not tie workers to their initial employer. This facilitates labour mobility and the flexibility of the labour market. It is

also consistent with the emancipation of the worker, who becomes more independent of specific employers. Finally, many social-insurance programs suffer from the problem that it is hard to separate the truly needy from other individuals who do not really need help from the government. If social norms regarding the take-up of benefits are endogenous and the take-up rate depends positively on how many people already receive benefits (as argued by Lindbeck (2006)), individual accounts may improve the sustainability of the welfare state by inducing people not to take up social benefits unless they really need them. This helps to halt an erosion of social norms. With individual accounts reducing moral hazard for middle- and higher incomes, the government can focus its active labor-market policies more on the life-time poor, thereby also protecting the social norms of this group.

Individual accounts also have implications that have not been included in our formal analysis. By separating lifetime redistribution from consumption-smoothing and insurance, individual accounts increase the transparency of lifetime redistribution. This may weaken the political support for this redistribution. Another factor that may work in the same direction is that the middle class no longer benefits from redistribution, which is now more closely targeted at the lifetime poor. At the same time, however, individual accounts give individuals a stronger sense of ownership and personal responsibility. This may strengthen popular support for the welfare state and the liquidity and lifetime insurance it provides. Stronger personal ownership may also make it more difficult for the government to change benefit rules, thereby reducing political risks.

The lifetime income insurance built into the system limits the cuts in effective social benefits to high- and middle incomes in order to contain the possible adverse effects on the incomes of the lifetime poor. This may encourage the middle and higher income earners to lobby for stronger employment protection, thereby harming the flexibility of the labor market. The lifetime income guarantee implies also that, while marginal rates are cut for others, marginal tax rates remain large only for the lifetime poor.¹⁶ The employment gap between low-skilled and high-skilled workers may thus increase unless

¹⁶A related drawback is that, although the lifetime poor may not become worse off in absolute terms, they may become poorer compared to the lifetime rich. This is a serious drawback if people care more about relative incomes than absolute incomes. In the presence of such standard-of-living utilities, optimal marginal tax rates at the top of the income distribution would be positive.

the government focuses active labour-market policies on the bottom of the labour market and employs instruments other than financial incentives to activate the lifetime poor. Hence, in contrast to those with higher lifetime incomes, these individuals may face more government intrusion in their private lives and are less free to make their own sovereign decisions.

If individuals lack the willpower or cognitive abilities to smooth consumption over their lifetimes, then annual disposable income becomes an important welfare indicator in addition to lifetime income. Accordingly, the government should base its redistributive policies not only on lifetime incomes (on the basis of the balances in individual accounts), but also on disposable incomes at each point in time. Intuitively, in the presence of myopia the government cannot rely on individuals to allocate their lifetime incomes optimally over their life course. In practice, while some consumers are myopic, others seem to be forward-looking, as mentioned earlier. This suggests that the optimal redistribution policy should be based on annual as well as lifetime incomes.

The analysis in this paper indicates that mandatory individual savings accounts can be a useful component of an overall social policy package. In addition to equal opportunities at the start of life through an equal distribution of human capital, such a policy package should provide some form of life-time income guarantee. By using information on lifetime incomes, redistribution implicit in such an income guarantee can occur at lower efficiency costs. Moreover, actuarially fair links between contributions and expected benefits alleviate the labor-market distortions associated with social insurance for middle- and high incomes. Finally, by facilitating consumption-smoothing through saving schemes offering liquidity insurance, the government increases the scope for self-insurance, thereby combating moral hazard in social insurance. Through all these channels, saving accounts support social policy by reducing the costs that are associated with an effective mix of redistribution, social insurance and consumption smoothing.

Technical appendix

A.1. The effects of individual savings accounts on the public budget

This section derives the formulas (4.1) and (4.2) that were used in section 5.3 to estimate the revenue effects of introducing IAs. We measure all variables in growth-adjusted present-value terms, assuming that the growth-adjusted real interest rate on government bonds is zero. Since individuals who end up with a deficit on their IA pay the same taxes and receive the same transfers as under the current tax-transfer system, we focus on those individuals who manage to accumulate a surplus on their IA at the date of retirement. For simplicity, we abstract from any changes in the revenue from capital income taxes stemming from changes in economic activity and in saving behavior.

Since real-world tax systems are piecewise linear, we assume a linear system of labour income taxation where the tax bill (T) of a person participating in the labour market is

$$T = \tau wh - I. \quad (\text{A.1})$$

Here τ is the marginal effective tax rate on labour income, including social security taxes as well as indirect taxes; w is the wage rate; h is the number of hours worked; and I is 'virtual' income, i.e., a parameter that may be calibrated to obtain a realistic value of the total average effective tax rate on labour income, given the various deductions and the form of the tax schedule imposed on intramarginal labor income.

Assuming a zero growth-adjusted real interest rate, and setting the total time available up until the official retirement age equal to unity, we can write the balance (A) on the IA at that time as

$$A = sewh - ab(1 - e), \quad 0 \leq \alpha \leq 1, \quad (\text{A.2})$$

where s is the rate of mandatory contribution to the IA; e is the average labour force participation rate over the active life of the representative wage earner with an IA surplus (so that ewh is his/her total labour income); b is the average after-tax public transfer received in periods of non-employment by people below the official retirement age, and α is the fraction of benefits to people of working age that is debited to the IAs.

Using (A.1) and (A.2), and assuming that the IA system is integrated in the public budget, we can write the growth-adjusted present value of the total net revenue (R)

collected from the representative member of a cohort with an IA surplus as¹⁷

$$\begin{aligned} R &= eT + sewh - (1 - e)b - P - A \\ &= e(\tau wh - I) - (1 - \alpha)(1 - e)b - P, \end{aligned} \quad (\text{A.3})$$

where P is the ordinary retirement benefit granted to people above the official retirement age. We see from the second line in (A.3) that the contribution rate s has no revenue effect, because all contributions are effectively remitted to individuals with a positive IA balance.

The introduction of IAs means that part of the labour income tax is replaced by a mandatory IA contribution and that part of the benefits received during periods of non-employment is debited to the IA. In formal terms, such a reform thus implies a cut in τ combined with a rise in s and α . We wish to estimate the revenue effect of introducing a system of IAs, starting from an initial situation without such a system where $A = s = \alpha = 0$. Using this initial condition and recalling that the proposed IA system does not involve any change in ordinary retirement benefits (i.e., $dP = 0$), we find from (A.3) that the revenue effect of introducing IAs amounts to

$$dR = \overbrace{ewh \cdot d\tau + b(1 - e) \cdot d\alpha}^{\text{mechanical effect}} + \overbrace{(T + b) \cdot de + \tau ew \cdot dh}^{\text{behavioral effect}}. \quad (\text{A.4})$$

¹⁷This specification assumes that contributions to the IA are not deductible; that IA balances are not taxed, and that only net (after-tax) benefits are debited to the IA. The DEC proposal described in section 4.1, in contrast, assumes that *pre-tax* benefits are debited to the IA and that IA contributions are deductible from the personal income tax base whereas IA balances are subject to personal income tax. In this case one can show that (A.3) modifies to

$$\begin{aligned} R &= e(\tau wh - I) - (1 - \alpha)(1 - e)b - P \\ &+ (m^A - m^w) sewh - \alpha(1 - e)b \left(\frac{m^A - t^b}{1 - t^b} \right), \end{aligned} \quad (\text{A.3.a})$$

where m^A is the marginal personal tax rate on IA balances, m^w is the marginal personal tax rate on labour income, and t^b is the average tax rate on benefit income. However, in the initial pre-reform equilibrium we have $s = \alpha = 0$, so to a first-order approximation, changes in e and h will have no impact on R via the last two terms on the right-hand side of (A.3.a). Hence, an analysis based on (A.3) approximates the revenue effect of the reform.

Note also that all variables in (A.3) are measured *after* indirect taxes, so the revenue effects of indirect consumption taxes are implicitly included (see the specification of effective tax rates in (A.8) below).

The so-called 'mechanical effect' indicated in (A.4) is the hypothetical effect on revenue that would materialize if the IA reform did not affect behavior. However, since the reform reduces the effective marginal and average tax rate on labour income as well as the effective benefit rate (because agents end up paying a fraction α of their own benefits via a reduced IA balance), it will affect labour force participation as well as the number of hours worked by those who participate, as also indicated in (A.4). For simplicity, we shall abstract from income effects on labour supply, since most recent empirical studies find that these effects are very small.¹⁸ Income effects will be absent if utility functions take the quasi-linear form

$$U = C - D \cdot [f(h) + q], \quad f' > 0, \quad f'' < 0,$$

where C is consumption, $f(h)$ is the disutility of working h hours, q is a fixed (pecuniary and/or psychological) cost of labour force participation, and D is a dummy variable taking the value of unity when the individual participates in the labour market and the value of zero when he/she does not participate. Following Immervoll et al. (2005), suppose q varies in a smooth continuous manner within a group of workers earning the same wage rate w . The participation rate of that group will then vary continuously with changes in the variable

$$y \equiv wh_o(1 - \tau) + I - b(1 - \alpha), \quad (\text{A.5})$$

representing the difference between net income when working and net income when not working, measured at the initial level of working hours, h_o . Note that a marginal change in h induced by a policy reform does not affect the utility of an employed worker, since the resulting change in consumption is offset by a change in the disutility of work when the initial working hours h_o have been optimized (i.e., $f'(h_o) dh = dC = w(1 - \tau) dh$ in the initial optimum). Hence, a change in h does not affect the incentive to participate in the labour market. This is why the variable y in (A.5) is measured at the given initial level of working hours.

Turning to labour supply at the intensive margin, in the absence of income effects the working hours of an employed worker depend exclusively on the marginal after-tax wage

¹⁸For example, using Danish data, Frederiksen et al. (2001) estimate an average income elasticity of around -0.005 .

rate, $w(1 - \tau)$. Defining the elasticities

$$\eta \equiv \frac{de/e}{dy/y}, \quad (\text{participation elasticity}),$$

$$\varepsilon \equiv \frac{dh/h}{dw(1 - \tau)/w(1 - \tau)}, \quad (\text{hours-of-work elasticity}),$$

we may write (A.4) as

$$dR = ewh \cdot d\tau + b(1 - e) \cdot d\alpha + \eta(T + b) \cdot \frac{edy}{y} + \varepsilon\tau ewh \cdot \frac{dw(1 - \tau)}{w(1 - \tau)}. \quad (\text{A.6})$$

Using

$$dw(1 - \tau) = -w \cdot d\tau,$$

$$y = wh_o - (T + b),$$

$$dy = -wh_o \cdot d\tau + b \cdot d\alpha,$$

and defining

$$t \equiv \frac{T}{wh}, \quad (\text{average labour income tax rate}),$$

$$c \equiv \frac{b}{wh}, \quad (\text{replacement rate}),$$

we can rewrite equation (A.6) as

$$\frac{dR}{ewh} = \left[1 - \varepsilon \left(\frac{\tau}{1 - \tau} \right) - \eta \left(\frac{t + c}{1 - t - c} \right) \right] d\tau + c \left[\frac{1 - e}{e} + \eta \left(\frac{t + c}{1 - t - c} \right) \right] d\alpha. \quad (\text{A.7})$$

Accounting for indirect taxes, we can write the effective marginal and average tax rates on labour income as

$$\tau = \frac{\tau^d + t^c}{1 + t^c}, \quad t = \frac{t^d + t^c}{1 + t^c}, \quad (\text{A.8})$$

where τ^d and t^d are the marginal and average direct tax rates, respectively, and t^c is the overall effective indirect tax rate on consumption. To apply formula (A.7), we also need to account for the link between $d\tau$ and $d\alpha$. We assume that the contribution to the IA is matched by a corresponding reduction in the marginal labour income tax rate so that $ds = -d\tau$. In the *absence* of changes in labour supply behavior, and given that we start out with $s_o = \alpha_o = A_o = 0$, the IA account balance for a person with a positive balance will then be

$$A = ewh \cdot ds - b(1 - e) \cdot d\alpha = -ewh \cdot d\tau - b(1 - e) \cdot d\alpha \quad \iff$$

$$d\alpha = - \left(\frac{ewh}{b(1-e)} \right) \left(d\tau + \frac{A}{ewh} \right). \quad (\text{A.9})$$

For the average Danish household with an IA surplus, the data used to produce the estimates presented in Table 6 imply that

$$\frac{A^p}{ewh} = 0.035, \quad \frac{ewh}{B(1-e)} = 9.38,$$

where A^p is the average account balance before tax, and B is the average rate of pre-tax benefit. In Denmark, transfer income is subject to the same tax schedule as labour income. Hence we assume that benefits and labour income are taxed at the same rate so that $b = B(1-t)$. Moreover, when applying formula (A.9), we assume (in accordance with the DEC proposal) that the account balance A^p is also taxed at the rate t , since contributions to the IA are deductible from the labour income tax base (see also footnote 17). In formula (A.9) we thus set $A/ewh = (1-t)A^p/ewh = (1-t)0.035$.

In a Danish context, plausible estimates for the effective tax rates are

$$t^d = 0.42, \quad \tau^d = 0.54, \quad t^c = 0.26.$$

The estimate for t^d is taken from the OECD Taxing Wages report (OECD (2005)) and refers to the average Danish production worker. The estimate for the average value of the marginal direct tax rate on labour income (τ^d) is taken from the Danish Ministry of Finance (2004), and the estimate for t^c is based on Carey and Rabesona (2004, Table 7.B2) and is an average figure for Denmark for the period 1990-2000.

We finally need to estimate $c \equiv b/wh = B(1-t)/wh$. In the Danish system of unemployment insurance, the gross replacement rate B/wh was recently estimated by the Confederation of Danish Trade Unions to be about 0.55 for the average production worker. However, in many other transfer programs the average replacement rate is somewhat lower, so we choose to set $B/wh = 0.5$.¹⁹ Using (A.8) and the estimates of t^d , τ^d and t^c , we then find

$$c = B(1-t)/wh = 0.23.$$

¹⁹Our estimated benefit rates b and B are averages across all individuals and social insurance programs, regardless of whether or not the program is included in the IA system. Important programs such as social assistance and ordinary retirement benefits have relatively low replacement rates. We are thus confident that our procedure does not overestimate the average replacement rate in the programs included in the IA system.

Given these parameter values and some assumptions on the labour supply elasticities, one may use the formulas (A.7) through (A.9) to estimate the revenue effect of the IA reform proposed by the Danish Economic Council. The resulting estimate is reported in section 5.3 of the main text.

A.2. Incentive and welfare effects of individual savings accounts for social insurance

This section presents a simple formal framework to illustrate how the introduction of individual accounts may enable the government to provide lifetime income insurance and liquidity insurance in a more efficient manner. Our framework is a simplified version of the more elaborate model developed in Bovenberg and Sørensen (2004).

Individuals in our model economy live for two periods. During the first period each employed person works full time, but a fraction of each young cohort is involuntarily unemployed due to negative labour market shocks. All unemployment risks are borne by young workers because jobs are rationed on a Last-In-First-Out basis. Young unemployed workers are liquidity-constrained and thus undertake no savings during their youth. In the second period of life, people choose to work only a fraction e of the time and are retired from the labour force during the remaining fraction $1 - e$ of that period. We assume that the retirement age chosen by the consumer is no lower than the age limit entitling people to (early) retirement benefits. For simplicity, we take pre-tax factor prices as given and abstract from taxes on capital income.

Household preferences and budget constraints

For all consumers in the economy, lifetime utility \widehat{U} is given by the utility function

$$\widehat{U} = U_1(C_1) + \left(\frac{1}{1+\delta}\right) [U_2(C_2) - h(e)] \quad (\text{A.10})$$

$$U'_i > 0, \quad U''_i < 0, \quad i = 1, 2; \quad h' > 0, \quad h'' > 0, \quad \delta > 0, \quad 0 \leq e \leq 1$$

where U_i is instantaneous utility from consumption, C_1 is first-period consumption, C_2 is second-period consumption, δ is the rate of time preference, e is the second-period employment rate (the fraction of the second period during which the consumer works before he retires), and $h(e)$ is the disutility from second-period work. Since the working hours of employed young workers are institutionally fixed, and since a young worker is either fully

employed or fully unemployed, the disutility from first-period work is exogenous and is therefore ignored in (A.10).

Households are divided into those who are exposed to involuntary unemployment during their youth and those who are not. The latter group will be called 'high-income earners', denoted by superscript h . A high-income earner, who has an employment rate of unity during the first period of his life, is subject to the budget constraints

$$C_1^h = w(1 - t - s) + y_1 - S, \quad S \geq 0, \quad (\text{A.11})$$

$$C_2^h = (1 + r)S + we^h(1 - t - s) + y_2(1 - e^h) + A^h, \quad (\text{A.12})$$

$$A^h = (1 + r)sw + swe^h - \alpha_2 y_2(1 - e^h), \quad A^h \geq 0, \quad 0 \leq \alpha_2 \leq 1, \quad (\text{A.13})$$

where w is the real wage rate before tax, t is the labour income tax rate, s is the mandatory rate of social security contribution to the consumer's individual account, y_1 is an in-work benefit available to young workers, S is financial saving (excluding the contribution to the IA), r is the real interest rate (determined in the world capital market), y_2 is an early retirement benefit, and A^h is the balance on the consumer's individual account that is paid out in the second period when he retires. From (A.13), we see that this balance consists of the contributions to the IA made during young age, with interest added, $(1 + r)sw$, plus the contribution made in the second period, swe , minus the fraction α_2 of the early retirement benefit $y_2(1 - e^h)$ that is financed by debiting his individual account. In the case of $\alpha_2 = s = 0$, we have a conventional tax-financed system of social insurance without IAs. When $\alpha_2 > 0$, part of the individual's (early) retirement benefit is financed by withdrawals from his IA. The constraint $A^h \geq 0$ reflects the lifetime income insurance built into the IA system: if the balance on the IA is negative at the time of retirement, the account is set at zero, and the individual still receives his ordinary retirement benefit. The constraint $S \geq 0$ indicates that the individual is unable to borrow against his expected future labour and transfer income. For high-income earners who are fully employed during their youth, we assume that the two constraints are not binding. We may then consolidate (A.11) through (A.13) into the single lifetime budget constraint

$$C_1^h + \frac{C_2^h}{1 + r} = w(1 - t) + y_1 + \frac{we^h(1 - t) + y_2(1 - \alpha_2)(1 - e^h)}{1 + r}. \quad (\text{A.14})$$

The social security tax rate s has dropped out of (A.14). This indicates that, for a consumer who is never liquidity-constrained, the mandatory contribution to the IA is

a perfect substitute for voluntary private saving and will hence leave his total saving $S + sw$ unchanged. Equation (A.14) also shows that, for an unconstrained consumer, the introduction of IAs (a positive value of α_2) amounts to a reduction in the effective rate of early retirement benefit, $y_2(1 - \alpha_2)$, in addition to a cut in marginal and effective tax rate t .

The representative high-income earner maximizes lifetime utility (A.10) subject to the budget constraint (A.14). The solution to this problem can be shown to yield an indirect lifetime utility function of the form $V^h = V^h(y_1, t, y_2, \alpha_2)$ with the derivatives

$$V_{y_1}^h \equiv \frac{\partial V^h}{\partial y_1} = \lambda^h, \quad V_t^h \equiv \frac{\partial V^h}{\partial t} = -\lambda^h \left(w + \frac{we^h}{1+r} \right), \quad (\text{A.15})$$

$$V_{y_2}^h \equiv \frac{\partial V^h}{\partial y_2} = \frac{\lambda^h (1 - \alpha_2) (1 - e^h)}{1+r}, \quad V_{\alpha_2}^h \equiv \frac{\partial V^h}{\partial \alpha_2} = -\frac{\lambda^h y_2 (1 - e^h)}{1+r}, \quad (\text{A.16})$$

where λ^h is the marginal utility of first-period income.

Those who are exposed to involuntary unemployment during the first period of life are divided into 'low-income earners' (denoted by superscript l) and 'medium-income earners' (indicated by superscript m). For both groups, the income loss from unemployment is so severe that the dissaving constraint $S \geq 0$ becomes binding, implying that no savings are made out of the net unemployment benefit b received during the first period. Hence

$$C_1^i = b, \quad i = l, m. \quad (\text{A.17})$$

In the second period, a medium-income earner earns the normal wage rate w , which is sufficiently high to enable him to accumulate a surplus on his IA. Thus a medium-income earner is subject to the second-period constraints

$$C_2^m = we^m (1 - t - s) + y_2 (1 - e^m) + A^m, \quad (\text{A.18})$$

$$A^m = swe^m - \alpha_b b (1 + r) - \alpha_2 y_2 (1 - e^m), \quad A^m \geq 0, \quad (\text{A.19})$$

where the term $-\alpha_b b (1 + r)$ on the right-hand side of (A.19) is the net balance on the IA carried over from the first to the second period (with interest added), given that a fraction α_b of the unemployment benefit received is debited to the IA. Inserting (A.19) into (A.18) to eliminate A^m , we obtain

$$C_2^m = we^m (1 - t) - \alpha_b b (1 + r) + y_2 (1 - \alpha_2) (1 - e^m). \quad (\text{A.20})$$

Equations (A.17) through (A.20) illustrate how the effect of IAs differs from the effect of a simple cut in the rate of unemployment benefit: a cut in b would force the worker to reduce his consumption in the first period of life, when the marginal utility of consumption is relatively high due to the liquidity constraint. By contrast, the introduction of individual unemployment accounts (an increase in α_b from zero to some positive number) would force a cut in consumption only in the second period when the marginal utility cost would be lower (since the consumer can escape liquidity constraints in the second period by postponing his date of retirement). This observation is one of the keys to understanding the potential for welfare gains from IAs. We also note from (A.17) and (A.20) that, as long as $A^m \geq 0$, the social security contribution s does not affect the behavior or welfare of the liquidity-constrained individual (since this contribution is essentially money that he pays to himself).

The medium-income earner maximizes (A.10) subject to (A.17) and (A.20). This yields an indirect utility function of the form $V^m = V^m(b, t, y_2, \alpha_b, \alpha_2)$ with derivatives

$$V_b^m \equiv \frac{\partial V^m}{\partial b} = \lambda_1^m - \frac{\lambda_2^m \alpha_b (1+r)}{1+\delta}, \quad V_t^m \equiv \frac{\partial V^m}{\partial t} = -\frac{\lambda_2^m w e^m}{1+\delta}, \quad (\text{A.21})$$

$$\begin{aligned} V_{y_2}^m &\equiv \frac{\partial V^m}{\partial y_2} = \frac{\lambda_2^m (1-\alpha_2)(1-e^m)}{1+\delta}, & V_{\alpha_b}^m &\equiv \frac{\partial V^m}{\partial \alpha_b} = -\frac{\lambda_2^m b(1+r)}{1+\delta}, \\ V_{\alpha_2}^m &\equiv \frac{\partial V^m}{\partial \alpha_2} = -\frac{\lambda_2^m y_2(1-e^m)}{1+\delta}, \end{aligned} \quad (\text{A.22})$$

where $\lambda_1^m \equiv U'_1(b)$ is the marginal utility of first-period income, and λ_2^m is the marginal utility of second-period income.

A low-income earner earns a wage rate θw ($\theta < 1$), which is so low that he ends up with a negative balance on his IA. Due to the lifetime income guarantee, his IA balance is therefore set at zero at the time of retirement, implying a second-period budget constraint

$$C_2^l = \theta w e^l (1-t-s) + y_2 (1-e^l), \quad (\text{A.23})$$

yielding an indirect utility function of the form $V^l = V^l(b, t, s)$ with the properties

$$V_b^l = \lambda_1^l \equiv U'_1(b), \quad V_t^l = V_s^l = -\frac{\lambda_2^l w e^l}{1+\delta}, \quad V_{y_2}^l = \frac{\lambda_2^l (1-e^l)}{1+\delta}. \quad (\text{A.24})$$

In this case, the mandatory social security contribution s works just like an ordinary labour income tax.

The government budget

The present value of the net taxes paid by a cohort over its life cycle is measured by that cohort's *generational account*. Treating the IA system as a part of the public budget, and using (A.13), we may write the generational account of a cohort of high-income earners (g^h) as

$$\begin{aligned} g^h &= (t+s)w - y_1 + \frac{(t+s)we^h - y_2(1-e^h) - A^h}{1+r} \\ &= tw - y_1 + \frac{twe^h - y_2(1-\alpha_2)(1-e^h)}{1+r}. \end{aligned} \quad (\text{A.25})$$

For a cohort of medium-income earners, we can use (A.19) to write the generational account as

$$\begin{aligned} g^m &= -b + \frac{(t+s)we^m - y_2(1-e^m) - A^m}{1+r} \\ &= -b(1-\alpha_b) + \frac{twe^m - y_2(1-\alpha_2)(1-e^m)}{1+r}. \end{aligned} \quad (\text{A.26})$$

Finally, for the low-income earners who do not manage to accumulate a surplus on their IAs, the generational account amounts to

$$g^l = -b + \frac{(t+s)we^l - y_2(1-e^l)}{1+r}. \quad (\text{A.27})$$

Below we will use (A.25) through (A.27) to analyze how the introduction of IAs will affect the present value of net government revenue.

Lifetime income insurance through individual retirement accounts

Via the benefits b and y_2 , the tax-transfer system provides lifetime income insurance to low-income earners by guaranteeing a minimum lifetime income. We will now show that, by introducing IAs, the government can provide such insurance in a more efficient manner. Under the IA system, lifetime income insurance is ensured via the provision that negative IA balances are simply cancelled at the time of retirement so that the net benefits received can never fall below those offered in the absence of IAs. To illustrate the implications of this feature of the IA system in the most transparent manner, we will now temporarily abstract from the group of medium-income earners – although our main result goes through also in an economy with three income groups (see Bovenberg and Sørensen (2004)).

To prove that the introduction of IAs has the potential to generate a Pareto improvement, we show that an IA reform designed to keep the utility of all individuals constant will surely improve the public budget. In that case, the government is obviously able to make everybody better off, for example by using the extra revenue to raise the universal retirement benefit y_2 .²⁰

Suppose that, starting from $s = 0$, the government introduces a mandatory contribution to an individual retirement account while at the same time reducing the labour income tax rate by a similar amount so that $ds + dt = 0$, $dt < 0$. Recalling from (A.24) that $V_t^l = V_s^l$, such a reform will keep the lifetime utility of low-income earners constant, and according to (A.27) will have no impact on their generational account. Suppose further that the fraction of retirement benefits that is required to be financed via the IA (α_2) is calibrated so as to keep the lifetime utility of high-income earners constant. Using (A.15) and (A.16), we find that this requires

$$V_t^h \cdot dt + V_{\alpha_2}^h \cdot d\alpha_2 = 0 \quad \implies \quad \frac{d\alpha_2}{dt} = - \left[\frac{w(1+r) + we^h}{y_2(1-e^h)} \right] \quad (\text{A.28})$$

Differentiating (A.25), inserting (A.28), and recalling that $\alpha_2 = 0$ initially, we find that the effect of this reform on the generational account of high-income earners is

$$dg^h = \left(\frac{tw + y_2}{1+r} \right) \left(\frac{we^h(2+r)}{y(1-e^h)} \right) \varepsilon_c^h \cdot (-dt), \quad (\text{A.29})$$

$$y \equiv w(1-t) - (1-\alpha_2)y_2, \quad \varepsilon_c^h \equiv \left(\frac{\partial e^h}{\partial y} \right)_c \cdot \frac{y}{e^h},$$

where ε_c^h is the *compensated* elasticity of a high-income earner's second-period labour supply with respect to the net reward to work (y), defined as the difference between the after-tax wage rate and the net rate of retirement benefit. Since this elasticity measures a pure substitution effect, ε_c^h is positive, and since we also have $dt < 0$, it follows from (A.29) that the IA reform will boost net government revenue by inducing high-income earners to postpone their retirement. Since the reform was designed to keep everybody's utility constant, it follows that the government can generate a Pareto improvement, say, by channeling the extra revenue back to consumers via an increase in y_2 .²¹ Effectively, the

²⁰A similar procedure for the analysis of the welfare effects of policy reforms was previously used by Kaplow (1996) and Bovenberg and Sørensen (2004).

²¹Of course, this will to some extent dampen the positive effect of the IA reform on labour supply.

IA reform enables the government to engineer an efficiency-enhancing cut in taxes and retirement benefits for high-income earners without affecting the incentives and welfare of low-income earners.

Lifetime income insurance through individual unemployment accounts

We will now show that the introduction of individual unemployment accounts will also enable the government to offer lifetime income insurance in a more efficient way. To illustrate the point, we must now include the medium-income earners in the analysis, along with high-income and low-income earners. We follow the same procedure as before, designing an IA reform that will keep everybody's utility constant and demonstrating that this reform will improve the public budget, thus allowing a Pareto improvement.

The introduction of unemployment accounts involves raising s and α_b from zero to some positive amounts while at the same time reducing t . To keep the utility and the generational account of low-income earners constant, we maintain the assumption that $ds + dt = 0$, $dt < 0$. The derivatives V_t^m and $V_{\alpha_b}^m$ in (A.21) and (A.22) indicate the effects of the changes in t and α_b on the welfare of medium-income earners. These expressions reveal that a constant utility for this group requires

$$\frac{d\alpha_b}{dt} = -\frac{we^m}{b(1+r)}. \quad (\text{A.29})$$

Finally, using the expressions for V_t^h and $V_{y_1}^h$ in (A.15) and (A.16), we keep the utility of high-income earners constant by adjusting the in-work benefit y_1 so that

$$\frac{dy_1}{dt} = w + \frac{we^h}{1+r} \quad (\text{A.30})$$

Differentiating (A.25) and (A.26) and using (A.29) and (A.30), we now find that

$$dg^i = \left(\frac{tw + y_2}{1+r} \right) \left(\frac{we^i}{y} \right) \varepsilon_c^i \cdot (-dt) > 0, \quad \varepsilon_c^i \equiv \left(\frac{\partial e^i}{\partial y} \right)_c \frac{y}{e^i} > 0, \quad i = m, h, \quad (\text{A.31})$$

where ε_c^i is the compensated second-period labour supply elasticity of group i . Hence, by effectively allowing a selective tax cut for medium- and high-income earners, the introduction of unemployment accounts generates a higher labour supply from these two groups, thereby boosting public net revenue and enabling the government to engineer a Pareto improvement.

Note that even though the unemployment accounts reduce the present value of the net unemployment benefits paid out to medium-income earners, the IA system does not

reduce the consumption possibilities of these individuals during the first period when they are liquidity-constrained. Indeed, we shall now show that unemployment accounts enable the government to provide liquidity insurance in a more efficient manner than through the conventional tax-transfer system.

Liquidity insurance through unemployment accounts

To demonstrate this result in the simplest possible manner, we abstract from the group of low-income earners, but once again we emphasize that the qualitative result carries over to an economy with three income groups, as shown in Bovenberg and Sørensen (2004).

Under a traditional tax-transfer system, the liquidity constraints of unemployed workers may be alleviated through a tax-financed rise in unemployment benefits. The welfare effect of such a reform may be measured by its impact on the public budget, assuming that the reform is designed to keep everyone's utility constant. For a medium-income earner, (A.21) reveals that a constant utility level requires

$$V_b^m \cdot db + V_t^m \cdot dt = 0 \quad \implies \quad \frac{db}{dt} = \frac{\lambda_2^m w e^m}{\lambda_1^m (1 + \delta)}. \quad (\text{A.32})$$

Notice that while (A.32) ensures that a medium-income earner's *lifetime* utility is kept constant, this is achieved by raising his first-period utility through a rise in b and lowering his second-period utility by raising t . With a binding first-period liquidity constraint, the medium-income earner will not want to compensate for the fall in second-period utility by shifting consumption from the first to the second period; the reform will thus induce an *uncompensated* second-period labour-supply response, which may be decomposed via the Slutsky equation

$$\frac{\partial e^m}{\partial t} = \left(\frac{\partial e^m}{\partial t} \right)_c - w e^m \left(\frac{\partial e^m}{\partial I} \right), \quad (\text{A.33})$$

where $\left(\frac{\partial e^m}{\partial t} \right)_c$ is the substitution effect and $-w e^m \left(\frac{\partial e^m}{\partial I} \right)$ is the income effect. From (A.26), (A.32) and (A.33) we find the following impact of the tax-financed rise in unemployment benefits on the medium-income earner's generational account,

$$dg^m = \left\{ \overbrace{\left[\frac{\lambda_1^m}{\lambda_2^m \frac{(1+r)}{1+\delta}} - 1 \right]}^{\text{liquidity insurance effect}} - \left(\frac{\lambda_1^m (1 + \delta)}{\lambda_2^m} \right) \left(\frac{tw + y_2}{1 + r} \right) \overbrace{\left(\frac{e^m \varepsilon_I^m}{I} + \frac{\varepsilon_c^m}{y} \right)}^{\text{labour supply effect}} \right\} db, \quad (\text{A.34})$$

$$\varepsilon_I^m \equiv \frac{\partial e^m}{\partial I} \frac{I}{e^m} < 0,$$

where ε_I^m is the income elasticity of labour supply which is negative under the assumption that leisure is a normal good. With a binding liquidity constraint in the first period, the magnitude $\lambda_1^m / \lambda_2^m \left(\frac{1+r}{1+\delta} \right)$ will be greater than one, so the positive term in the square bracket in (A.34) reflects the efficiency-enhancing effect of improved liquidity insurance via higher unemployment benefits. The higher labour income tax rate associated with higher benefits will have a negative impact on the medium-income earner's second-period labour supply, however, assuming that the substitution effect (reflected in the compensated labour supply elasticity ε_c^m) dominates the income effect (captured by ε_I^m). Hence, the overall net effect on the medium-earner's generational account is uncertain.

The high-income earner's utility is kept constant by adjusting the in-work benefit y_1 in accordance with (A.30). This means that the impact on the high-income earner's generational account will still be given by the expression in (A.31). However, since we now have $dt > 0$, the effect will be negative, reflecting the efficiency loss from the larger tax distortion to the labour supply of high-income earners (note that since high-income earners are not liquidity-constrained, the relevant labour supply elasticity in (A.31) is still the compensated elasticity, which is unambiguously positive).

Suppose now that, instead of being financed through higher taxes, the rise in the rate of unemployment benefit is financed by debiting a fraction of the benefit to an individual unemployment account, with the net balance being paid out when the consumer retires.²² Since high-income earners experience no unemployment, they will be unaffected by such a reform; their IA contributions will simply be returned to them with interest when they retire. Financing the rise in unemployment benefits via individual unemployment accounts thus sidesteps the negative impact on the high-income earner's labour supply occurring under tax finance.

Consider next the medium-income earners. From (A.21) and (A.22) it follows that a reform satisfying

$$V_b^m \cdot db + V_{\alpha_b}^m \cdot d\alpha_b = 0 \quad \implies \quad \frac{d\alpha_b}{db} = \frac{\lambda_1^m (1 + \delta)}{\lambda_2^m b (1 + r)} \quad (\text{A.35})$$

²²We assume that the social security contribution s is sufficiently high to ensure a positive IA balance at the date of retirement.

will keep the lifetime utility of this group constant, but the debiting of the fraction α_b of the unemployment benefit to the IA will have a positive income effect on second-period labour supply:

$$\frac{\partial e^m}{\partial \alpha_b} = -b(1+r) \cdot \frac{\partial e^m}{\partial I} > 0. \quad (\text{A.36})$$

Using (A.26), (A.35) and (A.36), we obtain

$$dg^m = \left\{ \overbrace{\left[\frac{\lambda_1^m}{\lambda_2^m \left(\frac{1+r}{1+\delta} \right)} - 1 \right]}^{\text{liquidity insurance effect}} - \left(\frac{\lambda_1^m (1+\delta)}{\lambda_2^m} \right) \left(\frac{tw + y_2}{1+r} \right) \overbrace{\left(\frac{e^m \varepsilon_I^m}{I} \right)}^{\text{labour supply effect}} \right\} db. \quad (\text{A.37})$$

Since $\varepsilon_I^m < 0$, we see that the impact on the medium-income earner's generational account is now unambiguously positive, enabling the government to create a Pareto improvement. Moreover, comparing (A.34) and (A.37), we observe that the IA-financed rise in unemployment benefits yields a *stronger* positive impact on g^m than the tax-financed rise in benefits, since finance via unemployment accounts avoids the tax distortion on labour supply. Recalling that $dg^h < 0$ under tax finance whereas $dg^h = 0$ under finance via IAs, we may therefore conclude that it is more efficient to offer liquidity insurance by financing a rise in unemployment benefits through unemployment accounts than through higher taxes.

REFERENCES

- Bovenberg, A.L. (2005). Balancing work and family life during the life course. *De Economist* 153.
- Bovenberg, A.L. and P.B. Sørensen (2004). Improving the equity-efficiency trade-off: Mandatory savings accounts for social insurance. *International Tax and Public Finance* 11, 507-529.
- Bovenberg, A.L. and P.B. Sørensen (2006). Optimal taxation and social insurance in a lifetime perspective. EPRU Working Paper 06-01, Economic Policy Research Unit, University of Copenhagen.
- Caliendo, M. and S. Kopeinig (2005). Some practical guidance for the implementation of propensity score matching. IZA Discussion Paper No. 1588, Institute for the Study of Labor, Bonn.
- Campbell, J.Y. and N.G. Mankiw (1991). The response of consumption to income. *European Economic Review* 35, 723-756.
- Carey, D. and J. Rabesona (2004). Tax ratios on labor and capital and on consumption. Chapter 7 in P.B. Sørensen (ed.), *Measuring the Tax Burden on Capital and Labor*, MIT Press.
- Coronado, J.L., D. Fullerton and T. Glass (2000). The progressivity of social security. *NBER Working Paper* 7520, Cambridge MA.
- Danish Economic Council (2004). *Dansk Økonomi, Efterår 2004* (The Danish Economy, Fall 2004). Copenhagen.
- Danish Economic Council (2005). *Dansk Økonomi, Forår 2005* (The Danish Economy, Spring 2005). Copenhagen.
- Danish Ministry of Finance (2002) *Fordeling og Incitamentter 2002* (Distribution and incentives 2002). Copenhagen.
- Danish Ministry of Finance (2004) *Fordeling og Incitamentter 2004* (Distribution and incentives 2004). Copenhagen.
- Falkingham, J. and A. Harding (1996). *Poverty Alleviation Versus Social Insurance Systems: A Comparison of Lifetime Redistribution*. Contributions to Economic Analysis, vol. 232, North Holland, Amsterdam.

- Feldstein, M. and D. Altman (1998). Unemployment insurance savings accounts. *NBER Working Paper* 6860, Cambridge MA.
- Fölster, S. (1997). Social insurance based on personal savings accounts: A possible reform for overburdened welfare states?. *European Economy* no. 4, 1997, 81-100.
- Fölster, S. (1999). Social insurance based on personal savings. *The Economic Record* 75, 5-18.
- Fölster, S., R. Gidehag, M. Orszag, and D. Snower (2002). Assessing welfare accounts. *CEPR Discussion Paper* No. 3479.
- Frederiksen, A., E.K. Graversen and N. Smith (2001). Overtime work, dual job holding and taxation. IZA Discussion Paper No. 323, Institute for the Study of Labor, Bonn.
- Gustman, A.L. and T.L. Steinmeier (2000). How effective is redistribution under the social security benefit formula? *NBER Working Paper* 7597, Cambridge MA.
- Hansen, M.I. (2005). Medborgerkonti - Velfærdsydelse finansieret af individuelle opsparingskonti (Citizen Accounts - Social insurance financed by individual savings accounts). Master thesis, Department of Economics, University of Copenhagen.
- Hussénius, J. and J. Selén (1994). *Skatter och socialförsäkringar över livscykeln - En simuleringsmodell* (Taxes and social insurance across the life cycle - A simulation model). Ds 1994: 86 (ESO), Swedish Ministry of Finance, Stockholm.
- Immervoll, H., H.J. Kleven, C.T. Kreiner and E. Saez (2005). Welfare reform in European countries: A micro-simulation analysis. *Economic Journal* (forthcoming).
- Kaplow, L. (1996). The optimal supply of public goods and the distortionary cost of taxation. *National Tax Journal* 49, 513-533.
- Le Maire, D. and C. Scheuer (2005). Determinants of labour force participation for selected groups with weak labour market attachment: A panel data analysis for Denmark. Working Paper, Department of Economics, University of Copenhagen.
- Leijnse, F., K. Goudswaard, J. Plantinga, and J.P. van der Toren (2004). A different attitude to social security. Life Course Policy Conference, SER, the Hague, June 3-4 2004.
- Liebman, J.B. (2001). Redistribution in the current U.S. social security system. *NBER Working Paper* 8625, Cambridge MA.
- Lindbeck, A. (2006). The welfare state - Background, achievements, problems. Forth-

coming in the new edition of Palgrave.

Nelissen, J.H.M. (1998). Annual versus lifetime income redistribution by social security. *Journal of Public Economics* 68, 223-249.

O'Donoghue, C. (2001). Redistribution in the Irish tax-benefit system. Ph.D. thesis, London School of Economics, London.

OECD (2005). *Taxing Wages - Country Tables - Year 2003*. Paris.

Orszag, J. and D. Snower (1997). Expanding the welfare system: A proposal for reform. *European Economy* no. 4, 1997, 101-118.

Orszag, J. and D. Snower (2002). From unemployment benefits to unemployment accounts. IZA Discussion Paper No. 532, Institute for the Study of Labor, Bonn.

Orszag, J., P. Orszag, D. Snower and J. Stiglitz (1999). The impact of individual accounts: Piecemeal versus comprehensive approaches. Paper presented at the Annual Bank Conference on Development Economics, World Bank, Washington D.C., 1999.

Pettersson, T. and T. Pettersson (2003). Fördelningen ur et livscykelperspektiv (Distribution from a life cycle perspective). Bilaga 9 till långtidsutredningen, Swedish Ministry of Finance, Stockholm.

Rele, H. ter (2005). Measuring lifetime redistribution in Dutch collective arrangements. CPB Document No. 79, CPB Netherlands Bureau for Economic Policy Analysis, the Hague.

Statistics Denmark (2005). *Statistiske Efterretninger 2005:27*. Copenhagen.

Stiglitz, J. and J. Yun (2002). Integration of unemployment insurance with retirement insurance. *NBER Working Paper* 9199.

Sørensen, P.B. (2003). Social insurance based on private savings accounts. Chapter 10 in S. Cnossen and H.-W. Sinn (eds.), *Public Finances and Public Policy in the New Millennium*, MIT Press.

Table 1. The distribution of annual income and lifetime income in Denmark (2002 income levels)¹

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	Average ²	Gini coefficient
Annual income (2002 cross section)												
Factor income	-2	11	36	62	85	103	122	143	172	269	100	0.438
Gross income ³	36	64	78	91	104	118	133	151	178	275	123	0.288
Disposable income	27	48	57	64	71	78	86	96	109	158	79	0.242
Lifetime income⁴												
Factor income	31	56	69	80	90	100	110	123	142	198	100	0.253
Gross income ³	77	93	102	110	117	124	133	144	161	214	128	0.160
Disposable income	56	64	69	74	77	81	85	91	99	125	82	0.127

1. Average factor incomes have been normalized to 100. All incomes are measured according to an equivalence scale allowing for economies of scale in household consumption.

2. Average income across the entire sample population.

3. Factor income plus pre-tax public transfers received.

4. Average income per year in the life cycle.

Source: Hansen (2005, Table 5.2).

Table 2. The distribution of social transfers across annual income deciles (percent of disposable annual income, 2002 cross section)

Transfer program	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	Average¹
Social assistance	15.8	8.5	2.5	0.9	0.4	0.2	0.1	0.1	0.0	0.0	1.4
Housing benefits	0.6	2.3	3.6	2.9	1.3	0.6	0.4	0.1	0.0	0.0	0.9
Disability benefits	1.9	4.3	6.3	6.4	6.9	3.7	2.8	1.5	0.9	0.3	3.0
Sickness benefits	2.1	2.7	2.9	2.5	1.8	1.4	1.0	0.7	0.4	0.2	1.2
Unemployment benefits	4.1	4.5	5.0	4.5	3.5	3.0	2.4	1.9	1.4	0.7	2.5
Child benefits	6.5	4.6	3.6	3.5	3.2	2.7	2.2	1.6	1.0	0.6	2.3
Education benefits	18.9	3.8	1.4	0.8	0.6	0.4	0.3	0.2	0.1	0.0	1.2
Parental leave benefits	0.8	0.6	0.5	0.5	0.4	0.3	0.2	0.1	0.1	0.0	0.3
Early retirement benefits	1.0	3.2	6.0	5.3	4.6	3.9	3.1	2.4	1.8	0.9	2.9
Ordinary retirement benefits ²	8.5	28.9	25.9	16.6	8.9	6.1	4.2	2.9	2.2	1.8	8.1

1. Average across the entire sample population.

2. Basic plus supplementary retirement benefits.

Source: Hansen (2005, Table 5.3)

Table 3. The distribution of social transfers across lifetime income deciles (percent of disposable lifetime income)

Transfer program	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	Average¹
Social assistance	6.9	2.4	1.4	1.0	0.8	0.6	0.5	0.4	0.4	0.3	1.2
Housing benefits	2.8	2.0	1.5	1.1	0.8	0.7	0.5	0.3	0.2	0.1	0.8
Disability benefits	12.2	8.2	5.8	3.8	2.9	2.3	1.7	1.4	1.0	0.6	3.3
Sickness benefits	1.9	1.7	1.4	1.3	1.1	0.9	0.8	0.7	0.5	0.3	1.0
Unemployment benefits	4.0	3.8	3.4	3.1	2.7	2.5	2.1	1.9	1.5	1.0	2.4
Child benefits	2.5	1.8	1.6	1.4	1.3	1.2	1.1	1.0	0.9	0.6	1.2
Education benefits	2.1	1.6	1.3	1.2	1.1	1.0	1.0	0.9	0.9	0.8	1.1
Parental leave benefits	0.2	0.2	0.2	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.1
Early retirement benefits	4.3	4.9	5.1	5.0	4.6	4.5	3.9	3.4	2.8	1.6	3.8
Ordinary retirement benefits ²	13.7	12.8	11.6	10.5	9.4	8.6	7.8	7.0	6.1	4.8	8.6

1. Average across the entire sample population.

2. Basic plus supplementary retirement benefits.

Source: Hansen (2005, Table 5.4).

Table 4. The normalized redistribution index for Danish transfer programs

Transfer program	Annual Income ¹	Lifetime Income ¹	Percentage share of total spending on social transfers (2004) ²
Social assistance	0.70	0.47	6.2
Housing benefits	0.35	0.39	4.4
Disability benefits	0.14	0.39	13.8
Supplementary retirement benefits	0.37	0.19	n.a.
Sickness benefits	0.19	0.18	8.3
Unemployment benefits	0.09	0.11	9.7
Child benefits	0.13	0.10	8.0
Education benefits	0.68	0.04	5.3
Early retirement benefits	0.00	0.04	10.8
Parental leave benefits	0.22	0.02	0.1
Basic retirement benefit	0.22	0.00	28.1 ³

1. Excess value of the redistribution index over the redistribution index for a uniform lump sum transfer.

2. The table excludes a number of minor programs accounting for 5.3 percent of total spending on social transfers.

3. Sum of basic and supplementary retirement benefits.

Sources: Hansen (2005, Tables 5.5 and 6.2) and Statistics Denmark (2005, Table 2).

Table 5. Interpersonal versus intrapersonal redistribution in Denmark (1,000 euros, 2002 income levels)¹

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	Average²
. Accumulated lifetime factor income	470	795	991	1146	1287	1421	1576	1755	2030	2859	1433
. Taxes 'reserved' for social insurance ($\sum_t T_{it}$)	211	277	311	338	363	388	418	454	512	711	398
. Transfers received over the life cycle ($\sum_t B_{it}$)	546	521	477	434	399	377	348	324	295	262	398
. Self-financed transfers received in the same year ($\sum_t SY_{it}$)	155	171	171	168	167	167	167	169	173	180	169
. Self-financed transfers received in another year (OY_i)	54	101	130	149	159	163	157	143	117	81	125
. For $N_i > 0$: Net transfers received over the life cycle ($\sum_t B_{it} - \sum_t SY_{it} = N_i$)	337	249	176	117	73	46	24	11	4	1	104
. For $N_i \leq 0$: Net taxes paid over the life cycle ($\sum_t T_{it} - \sum_t SY_{it} = -N_i$)	2	5	10	20	37	58	94	142	221	450	104

1. All magnitudes are total amounts accumulated over the life cycle, assuming a zero growth-adjusted real discount rate.

2. Average across the entire sample population.

Source: Hansen (2005, Table 5.7).

Table 6. Average payments to and from the individual accounts and account balances at the time of retirement across lifetime income deciles¹

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	Average
Lifetime income (index)	62	79	86	92	97	102	107	113	121	141	100
Accumulated payment into account in percent of accumulated withdrawal from account	34	56	72	84	97	109	123	141	161	210	100
After-tax account balance at retirement ² in percent of accumulated lifetime disposable income ³	0.1	0.4	0.7	0.9	1.2	1.4	1.8	2.2	2.5	3.3	1.6
Percent of adult population with positive account balance	7.2	17.1	27.7	36.3	43.0	51.2	57.2	65.8	71.0	79.7	45.6

1. The estimates assume a zero growth-adjusted real interest rate and unchanged behaviour.

2. Average account balance across the entire sample population, where negative account balances have been set to zero.

3. Accumulated income up until the official retirement age of 65; average across the entire sample population

Source: Hansen (2005, Table 6.4).