Network for Studies on Pensions, Aging and Retirement

Lans Bovenberg

Martin Hansen

Peter Sorensen

Individual Accounts and the Life-Cycle Approach to Social Insurance

Discussion Paper 2006 - 028

2006

December 2006

INDIVIDUAL ACCOUNTS AND THE LIFE-CYCLE APPROACH TO SOCIAL INSURANCE

A. Lans Bovenberg, CentER, Tilburg University,
Netspar, CEPR and CESifo*
Martin Ino Hansen, Statistics Denmark**
Peter Birch Sørensen, University of Copenhagen,
The Danish Economic Council, EPRU and CESifo***

Abstract: We explore the rationales for financing part of social insurance via mandatory individual 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 labour 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

*) CentER and Netspar, Tilburg University,
P.O. Box 90153, 5000 LE Tilburg, The Netherlands. E-mail: A.L.Bovenberg@uvt.nl
**) Statistics Denmark, Sejrøgade 11, 2100 Copenhagen Ø, Denmark.
E-mail: mih@dst.dk
***) Department of Economics, University of Copenhagen, Studiestraede 6,
1455 Copenhagen K, Denmark. E-mail: Peter.Birch.Sorensen@econ.ku.dk

APPENDIX 2 to paper on

"Individual accounts and the life-cycle approach to social insurance"

by A. Lans Bovenberg, Martin I. Hansen and Peter Birch Sørensen

ESTIMATING LIFETIME INCOMES

The empirical analysis in this paper employs the most recent estimates of lifetime incomes in Denmark, presented in the 2005 spring report of the Danish Economic Council (DEC, 2005). This appendix explains the methodology used by the DEC and evaluates the quality of the estimates.

Constructing synthetic life cycles: the matching procedure

As mentioned in section 2.1, the estimates of lifetime incomes are based on a comprehensive micro panel data set including a representative sample of 10 percent of the Danish population above the age of 18 and covering 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.

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. 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 socioeconomic 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, 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 age.

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. This 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).

While individuals are matched on the basis of their expected disposable income, 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. Further, 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. Finally, 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.

Adjusting for policy changes

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, the recorded actual tax and transfer payments are therefore replaced 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 rules in the period 2000-2002 prevail over the full life cycles of all individuals. This 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, the 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 - the procedure is unlikely to imply a systematic bias in the amount of taxes and transfers assigned to the various lifetime income deciles.

Quality control

A main concern regarding the construction of synthetic life cycles is the reliability of the matching procedure. To test the quality of the estimates obtained through the procedure described above, we have examined each of the critical transitions between any two years in a synthetic life cycle where data for two different individuals have been matched. In order to construct a complete life cycle one has to carry out a total of 9 matches. The remaining transitions (income changes) between two years in a life cycle are on the other hand estimated correctly since they represent the observed income changes for the same individual. Although the matching of different individuals is based on a large number (53) of socioeconomic characteristics, it is not possible to include all relevant characteristics since some of these are unobservable. Hence some level of uncertainty will occur in relation to the matching of different persons, so one would expect that the correlation between income in the current year and next year's income is lower in the years when a match between two different individuals occurs than in the years when data for the same individual are used.

To test the above hypothesis we have calculated the autocorrelation coefficients between the disposable incomes in years *t* and *t+1* in those years *t* in the synthetic life cycles where a match between different individuals takes place. As a comparison, we have also calculated the autocorrelation coefficients between actual income during equivalent transitions of individuals found in the data registers in both 2001 and 2002. In other words, if a match between different individuals occurs at age 50 in the constructed life cycles, we compare the income changes between ages 50 and 51 in these synthetic life cycles to the observed income changes for an 'actual' person who was 50 in 2001 and 51 in 2002

Table A.1 reports our findings. On average the autocorrelation of income is reduced by 29 percent (from 0.75 to 0.53) in the transition years when two different individuals are matched. The lower autocorrelation when two different individuals are matched suggests that our method for constructing incomes over the life course tends to understate the persistence of income shocks, thereby understating the differences in life-time incomes across various agents. However, since any full life cycle consists of 72 annual transitions of which only 9 involve a match, this bias towards lower persistence does not seem disturbing. Even large errors in the critical transitions will only have a small influence on the total variance in the lifetime income.

(Table A.1 about here)

As a further quality check we have also examined whether various important socioeconomic characteristics vary in a credible manner over the constructed life cycles. Specifically, we have compared a range of characteristics (education, fertility, marital status, sex and disposable income) in the synthetic life cycles with equivalent variables from cross sectional data. The average values of these variables by age are reported in figure A.1. Since the lifetime income model reflects the cohorts aged 43-50 years in 2002, the values of the socioeconomic variables in this age group observed in the cross sectional data ought to be closely related to the corresponding variables in the constructed life cycles. At the same time, when one moves further away from these cohorts, a larger deviation between the modelled individuals and the cross sectional data can be expected.

Panel a) in figure A.1 reports the educational level by age in the synthetic life cycles and in the Danish population in 2002. Danes who were 43-50 years of age in 2002 have a higher level of education compared to the older generations in the same year. In the older age groups the ‘synthetic’ individuals in the constructed life cycles therefore have a better education than reflected in the cross sectional data. On the other hand, individuals in the constructed life cycles have a lower level of education in the younger age groups (age 25-40) compared to the Danish population in the cross sectional data. These findings accord with expectations.

The number of children in the life cycles and in the cross section data is almost identical. Thus the fertility of the life cycle individuals closely resembles the fertility of the Danish population in 2002. Panel b) does suggest that the modelled individuals have their children at a younger age. Also, in the cross sectional data there are apparently more singles in the younger age groups which also seems reasonable (panel c). Finally, the share of females in the constructed life cycles is closely related to the cross sectional data (see panel d). As a result of the higher female life expectancy, this share increases dramatically past the age of 77. The increase in the share of females in the constructed life cycles is not quite as dramatic since the middle aged males of 2002 have an improved life expectancy.[1] All of the above findings are as expected.

The disposable income in panel e) is influenced by a wide range of variables. During the early years of the modelled individuals’ lives one would expect a higher income compared to the cross sectional data since the 43-50 year olds spent fewer years educating themselves. Later on in life when people have children and the income has to be shared by more individuals, the disposable income should fall (due to the equivalence scale). Since the modelled individuals have their children earlier than reflected in the cross sectional data, this drop in income should appear earlier in life. Finally, since the life cycle individuals are better educated than the equivalent Danish cohort of 2002, they should experience higher income during their older years. The above expectations conform to our findings (cf. panel e).

---

[1] This is apparent when studying the most recent life expectancy tables of the Danish population (Statistical Yearbook 2006, Table 31, Statistics Denmark).

(Figure A.1 about here)

In summary, the above analysis indicates that the quality of the matching procedure is satisfactory and that the constructed life cycles display a plausible time profile of the most important socioeconomic variables.

Table A.1: Autocorrelation coefficients between incomes in the critical transitions

Age	Same individuals 2001-02	Modelled individuals
18-19 years	0,64	0,53
26-27 years	0,72	0,62
34-35 years	0,89	0,49
42-43 years	0,60	0,55
50-51 years	0,76	0,46
58-59 years	0,67	0,41
66-67 years	0,84	0,62
74-75 years	0,88	0,59
82-83 years	0,76	0,47
Average	0,75	0,53

Source: Hansen (2005), Appendix 5.4

Figure A.1: The constructed life cycles compared to cross sectional data (2002)

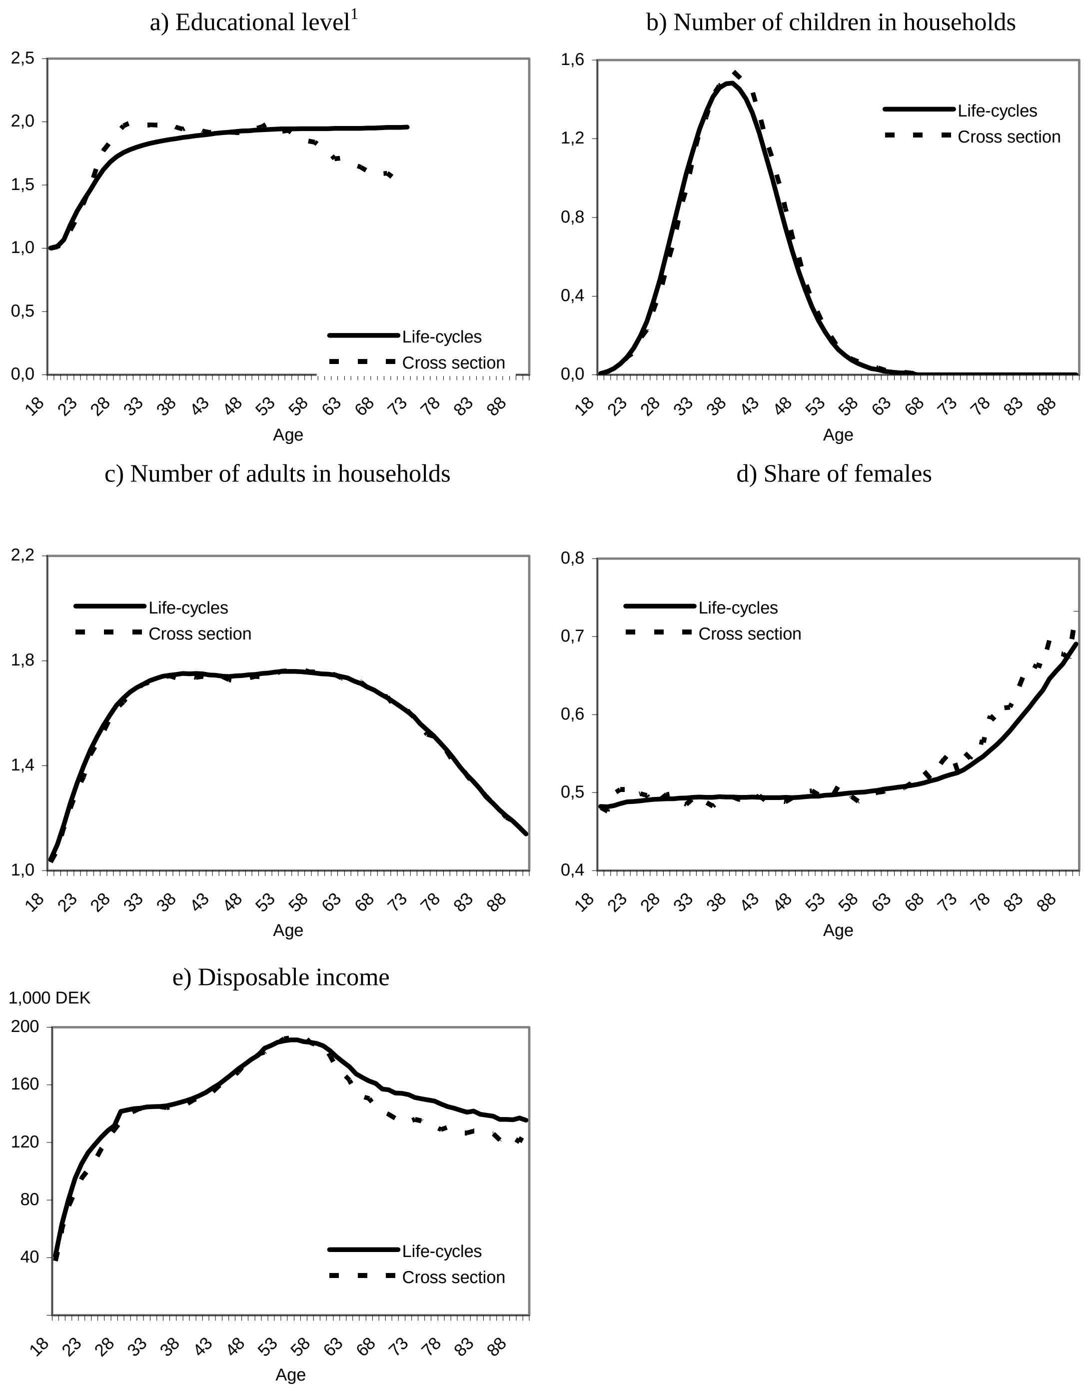

1) Individuals with no formal education (ufaglærte) are given the value 1. Medium skilled individuals (faglærte) are given the value 2 while high skilled individuals (videregående udd.) are given the value 3.

No information exists regarding the education of people born before 1923. Since the construction of the synthetic life cycles involves the matching of individuals in 2002 with other individuals of the same age in 1994, there is no information regarding the modelled individuals' educational level past the age of 72.

Source: Hansen (2005), Appendix 5.5.1 – 5.5.5

References

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 labour, Bonn.

Hansen, M.I. (2005). Medborgerkonti - Velfærdsydelser finansieret af individuelle opsparingskonti (Citizen Accounts - Social insurance financed by individual savings accounts). Master thesis, Department of Economics, University of Copenhagen.

APPENDIX 1 to paper on

"Individual accounts and the life-cycle approach to social insurance"

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

THE EFFECTS OF INDIVIDUAL ACCOUNTS ON LABOUR SUPPLY AND NET PUBLIC REVENUE

This appendix shows how to derive the formula presented in Box 4 and used in section 4 to estimate the effects on the public budget of introducing individual accounts.

The introduction of IAs involves a cut in τ combined with a rise in the variables s, α_e and α_m from zero to some positive numbers. Using equations (4.2) and (4.3) in Box 4;[1] remembering that $s = \alpha_e = \alpha_m = 0$ initially; and recalling that the proposed IA system does not involve any change in ordinary retirement benefits (i.e., $dy = 0$), we find that the revenue effect of introducing IAs amounts to

$$dR = \overbrace{ewh \cdot d\tau + b_e (1-e) \cdot d\alpha_e + b_m m \cdot d\alpha_m}^{\text{static effect}} + \overbrace{(\phi_e T + b_e) \cdot de - b_m \cdot dm + \phi_h \tau ew \cdot dh}^{\text{dynamic effect}} \tag{A.1}$$

The 'static effect' is the hypothetical effect on revenue that would materialize if the IA reform did not affect behaviour. However, since the reform reduces the effective marginal and average tax rate on labor income as well as the effective benefit rate, it will affect overall employment. The revenue implications of larger employment are given by the so-called 'dynamic effect' in (A.1), which is found by multiplying behavioural

[1] Equation (4.2) 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. Alternatively, one may assume 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 tax. In this case, by using $T = \tau wh - I$, one can show that (4.2) modifies to

$$\begin{aligned} R &= eT - (1-\alpha_e)(1-e) b_e - (1-\alpha_m) b_m m - y \\ &+ \left(\tau^A - \tau\right) sewh + \left(\frac{t^b - \tau^A}{1-t^b}\right) \left[\alpha_e (1-e) b_e + \alpha_m b_m m\right], \end{aligned} \tag{i}$$

where τ^A is the average tax rate on IA balances, and t^b is the average tax rate on benefit income. However, in the initial pre-reform equilibrium we have $s = \alpha_e = \alpha_m = 0$, 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 (i). Hence, an analysis based on (4.2) still approximates the revenue effect of the reform.

effects (de, dh and dm) by effective tax wedges. The parameter ϕ_e indicates to what extent the additional tax payments generated by additional labour force participation yields additional benefit rights. If the tax payments are actuarially fair, we have $\phi_e = 0$. Similarly, the parameter ϕ_h indicates to what extent the additional tax payments generated by additional hours worked yields entitles the taxpayer to additional benefits.

In modelling behavioural impacts, we abstract from income effects on labour supply, since most recent empirical studies find that these effects are quite small. Income effects will be absent if utility functions take the quasi-linear form

$$U = C - D \cdot [f(h) + q], \qquad f' > 0, \quad f'' > 0, \tag{A.2}$$

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 does not participate. Following Immervoll et al. (2007), 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

$$y \equiv wh_o(1 - \phi_e \tau) + \phi_e I - b_e(1 - \alpha_e), \tag{A.3}$$

representing the difference between net income when working and net income when not working, measured at the initial level of working hours, h_o. 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., $f'(h_o)\,dh = dC = w(1 - \phi_h \tau)\,dh$ 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.3) is measured at the given initial level of working hours. At the intensive margin of labour supply, the absence of income effects means that the working hours of an employed worker depend exclusively on the marginal after-tax wage rate, $w(1 - \phi_h \tau)$. Hence we define the labour supply elasticities

$$\eta \equiv \frac{de/e}{dy/y}, \qquad \text{(participation elasticity)}, \tag{A.4a}$$

$$\varepsilon \equiv \frac{dh/h}{dw(1-\phi_h\tau)/w(1-\phi_h\tau)}, \qquad \text{(hours-of-work elasticity)}. \tag{A.4b}$$

Moreover, the number of persons applying for a certain benefit (and hence our variable m) may depend on the benefit level, and employment may in part be affected by the variable m (e.g., the employment rate may depend on the number of people collecting education benefits). We therefore also define the elasticities

$$\chi \equiv \frac{dm/m}{db_m(1-\alpha_m)/b_m(1-\alpha_m)}, \qquad \text{(benefit dependency elasticity)}, \tag{A.4c}$$

$$\varphi \equiv -\frac{de}{dm}\frac{m}{e}, \tag{A.4d}$$

where χ could reflect a moral hazard effect. Using the elasticities in (A.4), we may write (A.1) as

$$\begin{aligned} dR &= ewh \cdot d\tau + b_e\,(1-e) \cdot d\alpha_e + b_m m \cdot d\alpha_m \\ &+ (\phi_e T + b_e)\left(\eta e \cdot \frac{dy}{y} - \varphi\chi e \cdot \frac{db_m(1-\alpha_m)}{b_m(1-\alpha_m)}\right) \\ &-\chi b_m m \cdot \frac{db_m(1-\alpha_m)}{b_m(1-\alpha_m)} + \varepsilon\phi_h \tau ewh \cdot \frac{dw\,(1-\phi_h\tau)}{w\,(1-\phi_h\tau)}. \end{aligned} \tag{A.5}$$

Defining

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

$$c_e \equiv \frac{b_e}{wh_o}, \qquad \text{(replacement rate)},$$

$$c_m \equiv \frac{b_m}{wh_o}, \qquad \text{(relative benefit rate)},$$

$$\bar{\phi}_h \equiv \frac{d(\phi_h\tau)}{d\tau}, \qquad \bar{\phi}_e \equiv \frac{d(\phi_e t)}{d\tau},$$

and using

$$dw\,(1-\phi_h\tau) = -w \cdot d(\phi_h\tau) = -w\bar{\phi}_h d\tau,$$

$$y = wh_o - (\phi_e T + b_e)\,,$$

$$dy = -wh_o\bar{\phi}_e \cdot d\tau + b_e \cdot d\alpha_e,$$

we can rewrite equation (A.5) as equation (4.4) in Box 4.

Reference

Immervoll, H., H.J. Kleven, C.T. Kreiner and E. Saez (2007). Welfare reform in European countries: A micro-simulation analysis. *Economic Journal* (forthcoming).

Attachments to

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

December 2006

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

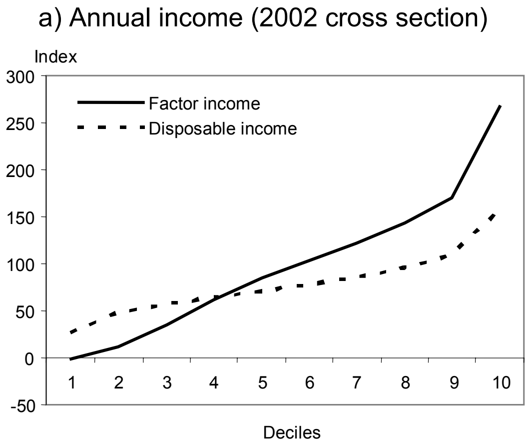

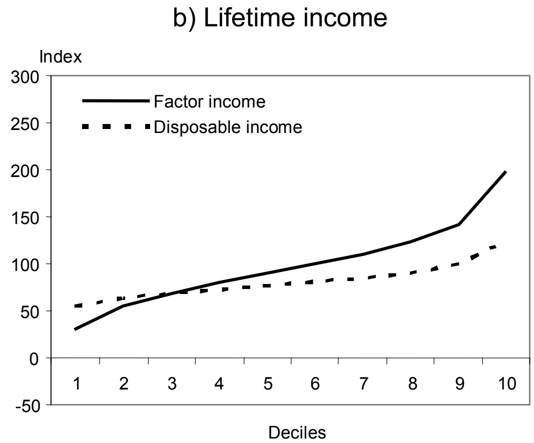

.

Figure 2. Nominal income growth and the after-tax nominal interest rate in Denmark[1]

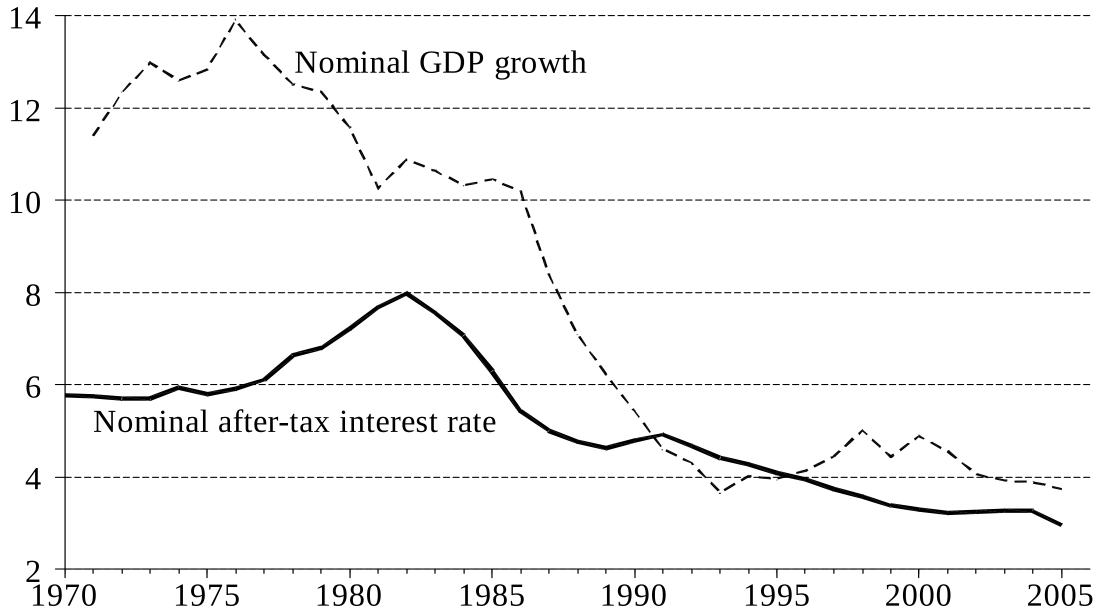

Notes:

1. The interest rate is the rate on 10-year Danish government bonds. The tax rate is a ‘typical’ tax rate on capital income. All numbers are five-year moving averages showing the average of the current and the four previous years.

Source: Statistics Denmark.

Table 1. The normalized redistribution index[1] for Danish transfer programs

Transfer program	Annual Income	Lifetime Income	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 insurance benefits	0.09	0.11	9.7
Child benefits	0.13	0.10	8.0
Grants to students in higher education	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. Total spending on social transfers amounted to 18.5 percent of GDP in 2004.
3. Sum of basic and supplementary retirement benefits.

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

Table 2: Comparison of individual accounts with alternatives

	Voluntary saving	Bismarckian insurance	Beveridgean redistribution	Individual accounts
Liquidity insurance	-	+	+	+
Lifetime redistribution	-	-	+	+
Low transaction costs	-	+	+	+
Paternalism protecting myopic individuals	-	+	+	+
Actuarial link between benefits and contributions	+	+	-	+/- [1)]
Self insurance limiting moral hazard	+	-	-	+/- [2)]

1) + for middle and high incomes; - for lifetime poor.
2) + for middle and high incomes if accounts are not used to buy insurance; - for lifetime poor

Table 3. Distributional characteristics of Danish transfer programs

Type of transfer	$\alpha_e c_e(1-e)/e$ and $\alpha_m c_m m/e$	$\alpha^*_e c^*_e(1-e^*)/e^*$ and $\alpha^*_m c^*_m m^*/e^*$	α_e and α_m	Redistribution index[1]
1. Early retirement benefits	0.0209	0.0292	0.4106	0.04
2. Sickness benefits[2]	0.0031	0.0081	0.0609	0.18
3. Parental leave benefits	0.0009	0.0015	0.0177	0.02
4. Short-term unemployment benefits[3]	0.0093	0.0212	0.1827	0.18
5. Long-term unemployment benefits[3]	0.0028	0.0083	0.0550	0.40
6. Disability benefits	0.0139	0.0255	0.2731	0.39
7. Total out-of-work benefits (1.+2.+3.+4.+5.+6.)	0.0509	0.0938	1.0	
8. Child benefits	0.0087	0.0128	0.5724	0.10
9. Grants to students in higher education	0.0054	0.0089	0.3553	0.04
10. Housing benefits	0.0011	0.0025	0.0724	0.39
11. Total of 8.+9.+10.	0.0152	0.0242	1.0	

1. Normalized redistribution index taken from Table 1.
2. All sickness benefits, including long-term benefits.
3. Unemployment insurance benefits plus social assistance related to unemployment.

Source: Own calculations, based on the lifetime incomes and transfers estimated in section 2.

Table 4. Average payments to and from the individual accounts and account balances at the time of retirement across lifetime income deciles[1]

	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).

Table 5. Estimated effects on net public revenue of including various transfer programs in the system of individual accounts[1]

Benefit	1. Static effect	2. Hours effect of lower taxes	3. Participation effect of lower taxes	4. Participation effect of lower benefits	5. Total dynamic effect (2.+3.+4.)	6. Total effect on net revenue (1.+5.)
Unemployment[2]	-1.19	0.18	0.57	1.00	1.75	0.56
Early retirement	-0.83	0.24	0.79	2.26	3.29	2.46
Sickness	-0.50	0.07	0.22	0.33	0.62	0.12
Parental leave	-0.06	0.01	0.04	0.10	0.15	0.09
All 4 programs above	-2.58	0.50	1.62	3.69	5.81	3.23
Education[3]	-0.35	0.07	0.24	-	-	-
Child	-0.41	0.11	0.34	-	-	-

Notes:
1. Measured in percent of the total labour income tax base for individuals with an IA surplus.
2. Only short-term unemployment benefits.
3. Only benefits to students in higher education.

Assumptions on parameter values: ε=0.05, η=0.10, τ=0.635, t=0.54, c_e=0.2, $\phi_h = \bar{\phi}_h = 0.987$, $\phi_e = \bar{\phi}_e = 0.985$

Source: Own calculations, based on Table 3 and formula (4.6) in Box 4.

INDIVIDUAL ACCOUNTS AND THE LIFE-CYCLE APPROACH TO SOCIAL INSURANCE

A. Lans Bovenberg, Martin Ino Hansen and Peter Birch Sørensen[1]

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 organisation 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 labour 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 therefore more relevant than ever before.

Against this background, we analyse 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 replace taxes that are currently financing social-insurance benefits. At

[1] We are grateful to Steen Jørgensen and Anne Kristine Høj for methodological advice regarding the estimation of lifetime incomes. We are also indebted to three anonymous referees, Guiseppe Bertola, Christian Gollier, Jukka Lassila and the Economic Policy panel for helpful comments on earlier drafts. Any remaining shortcomings are our own responsibility.

retirement, the remaining balances in 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 income 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 provide an empirical basis for our proposal for a social insurance reform, we start by estimating the degree to which current social-insurance programs succeed in redistributing lifetime incomes. We document that most social insurance transfers are much less redistributive when income is measured on a lifetime basis than when it is measured on an annual basis. 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 an ex-ante perspective (i.e. in terms of expected utilities before people know which shocks they are going to encounter during their lives). 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 labour-supply elasticities, the proposed reform would imply a Pareto improvement not only from an ex-ante but also from an ex-post point of view (i.e. in terms of realized utility in all contingencies) and would involve only a minor increase in the inequality of lifetime income distribution.

Fölster (1997, 1999), Orszag and Snower (1997, 2002), Feldstein and Altman (1998), Brown et al. (2006), Fölster et al. (2002), Stiglitz and Yun (2002), Sørensen (2003) and Bovenberg and Sørensen (2004) analyse 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, but in addition, we extend the existing literature in two directions. First, we quantify how the system of individual accounts affects work incentives, allowing us to compute the efficiency gains and the consequences for the government budget. Second, we design our proposal for individual accounts in such a way that it is Pareto improving.

Some of the U.S. literature addresses the effects of the individual accounts on saving behaviour, whereas we focus on the impact on labour-market behaviour. Indeed, while low saving rates are a prime policy concern in the United States, labour-market distortions depressing employment levels are a major focus of policy debate in the aging European welfare states. The main objective of the individual accounts is thus to boost labour supply rather than saving. Indeed, the accounts would be gradually phased in without intergenerational redistributional effects. The key property of our accounts is that they base taxation of labour income and the provision of social insurance benefits on lifetime income rather than annual incomes. By thus introducing a new instrument for redistribution that utilizes additional information on who is lifetime poor as a result of frequent adverse shocks over the life course, the accounts alleviate the labour-market distortions associated with interpersonal redistribution. At the same time, the positive labour-market effects of our policy proposal require that consumers are forward-looking, seeking to optimize their consumption flows over time. Hence, our theoretical analysis of the incentive and employment effects of individual accounts is cast in an explicitly intertemporal context (see Boxes 2 and 3 in section 3). The individual saving accounts 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 employs this information to offer a lifetime income guarantee (in the sense of bailing out individuals who end up with a negative account balance). Intuitively, the government can provide more efficient insurance of human capital risks by conditioning deductibles combatting moral hazard on aggregate income losses during the life course instead of having separate deductibles based on income losses in a each period. In this way, it provides better protection in the worst-case scenario of a succession of adverse shocks during the life course in exchange for less protection in other cases (see Gollier and Schlesinger (1995)).

A large literature exists on whether or not a transition from a pay-as-you-go to a funded pension system can result in a genuine Pareto improvement. Sinn (2000) and Belan and Pestieau (1999) conclude that the Pareto-improving transitions that have been found in some of this literature typically originate in the reduction of distortions that do not involve the funded nature of the new pension system as such. Hence, Paretian efficiency gains would materialize also if the government would alleviate the distortions without funding the pension system. Our paper arrives at a similar conclusion for social insurance benefits during the working life: the scope for Pareto improvement originates in the introduction of a more efficient redistributive instrument based on lifetime incomes rather than the possible funded nature of individual accounts.

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 labour-market incentives without harming the lifetime poor. Section 4 analyzes how individual accounts affect the distribution of lifetime incomes, the labour 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 finances public transfers that are channelled back to the same taxpayer at some point in the life cycle. This section uses Danish data to estimate 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 approach to social insurance laid out in section 3 and helps to identify those types of social insurance benefits that seem particularly well suited for finance through individual accounts. The present section also provides the data needed to quantify the effects of the reform proposal 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. The empirical analysis in this paper employs the most recent estimates of lifetime incomes in Denmark, presented in the spring 2005 report of the Danish Economic Council (DEC) (Danish Economic Council, 2005).[2] Box 1 explains the methodology used by the DEC; further details and documentation are provided in a technical appendix available at www.econ.ku.dk/pbs/).

The estimates of lifetime incomes are based on a comprehensive micro panel dataset covering the period 1994-2002 and comprising a representative sample of 10 percent of the

[2] Established 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.

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 income growth), using an equivalence scale adjusting for economies of scale in the consumption of multi-person households. While all monetary government transfers are included, the dataset does not include transfers in kind, i.e., the value of public services provided free of charge (or at subsidized prices) to 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 at the same point in the life cycle. This procedure implicitly abstracts from cohort effects other than those stemming from the 53 observable socioeconomic characteristics included in the dataset. By combining the observations for different cohorts, one obtains synthetic life cycles for individuals with different socioeconomic characteristics 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]

(Start of Box 1 about here)

Box 1. Estimating lifetime incomes

The idea underlying the estimation of lifetime incomes is the following. For the individuals included in the dataset, the evolution of actual incomes is known only for a nine-year interval (the period 1994-2002, both years included). To estimate a time series of annual incomes covering the whole (adult) life span of an individual with certain characteristics, it is therefore necessary to combine the data for different individuals who belonged to different age intervals during the period 1994-2002. For example, suppose person A was 18 years old in 1994 and was thus 26 years of age in 2002. Furthermore, suppose person B featured the same socioeconomic characteristics as A at the same age (e.g., in terms of gender, education etc.), except that he was 26 years old in 1994 and hence 34 years old in 2002. Finally, suppose that, at the age of 26,

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

both these persons had roughly the same income (measured in 2002 income levels). One may then consider A and B to be the same type of person, with A and B just representing different age intervals in the life cycle of their common type. Hence, one may combine the observed incomes earned by A during the age interval 18-26 with the observed incomes of B in the age interval 27-34 to obtain a series of incomes for the age interval 18-34 years in the life cycle of a person with the common characteristics of A and B. In a similar way, if person C features the same characteristics as A and B, was 34 years of age in 1994 and earned roughly the same income as B at that age, the observed incomes of person C during the data period 1994-2002 may be taken to represent the age interval 34-42 in the common life cycle of A, B and C. In the same way one may proceed to construct an income series covering the full life cycle of a person with certain characteristics at particular ages.

More precisely, the lifetime incomes were estimated using the following procedure. A person belonging to a group with certain characteristics who had age X in 2002 is matched with a person who was X years of age in 1994, but who had otherwise similar characteristics at the same age. 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.

First, all individuals of a given age are divided into 60 different groups, categorized according to gender, three different levels of education, and ten deciles of annual disposable income. In the second step, a person of a given age X belonging to a group with certain characteristics in 2002 is matched with another person who had a similar age and similar characteristics in 1994 and who featured an *expected* annual disposable income at age X that is as close as possible to the expected income of the first person. The expected disposable income is estimated by running several OLS regressions incorporating 53 different socioeconomic characteristics as explanatory variables, including family composition, detailed level of education, employment status, ethnic background etc. Matching individuals on the basis of expected rather than actual incomes allows the matching to exploit information on all the observable characteristics that tend to make the incomes of any two individuals converge. 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 at the same age.

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 different age groups during the years 2000-2002.

For simplicity, we assume that the relevant discount rate equals the average growth rate of real income (as documented in Figure X below, the average after-tax interest rate on long-term government bonds has in fact been quite close to the rate of wage growth in Denmark in recent years). One may then simply add up the annual incomes earned in each constructed life cycle to obtain an estimate of lifetime incomes.

(End of Box 1 here)

Figure 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 all numbers in the diagram 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, we have divided all lifetime incomes by the number of years in the constructed life cycle so as to obtain a measure of the average income earned per year. Disposable income is defined as factor income plus pre-tax public transfers received minus direct taxes paid. Since annual income is hump-shaped and varies considerably over the life cycle, one would expect the distribution of annual income to be much more unequal than the distribution of lifetime income. Figure 1 confirms this expectation. The figure indicates also that, in absolute terms (measured by the vertical distance between the solid and the dotted curves), the tax-transfer system has a stronger redistributive impact in an annual than in a lifetime perspective.[4]

(Figure 1 about here)

A main concern regarding the construction of synthetic life cycles is the reliability of the matching procedure. To test the quality of the estimates obtained through the procedure described in Box 1, we have examined each of the critical transitions between any two years in a synthetic life cycle where data for two different individuals have been matched. In order to construct a complete life cycle on the basis of our dataset, we had to carry out a total of nine matches. The remaining transitions (income changes) between two years in a life cycle are estimated correctly since they represent the observed income changes for the same individual. Although the matching of different individuals is based on a large number (53) of socioeconomic characteristics, it is not possible to include all relevant characteristics since some of these are unobservable. Some level of uncertainty thus occurs in relation to the matching of different persons. One would therefore expect that the correlation between income in the current year and next year's income is lower in the years when a match between two different individuals occurs than in the years when data for the same individual are used.

The technical appendix available at www.econ.ku.dk/pbs/ indicates that our method for constructing synthetic life cycles does indeed tend to understate somewhat the persistence of income shocks (and therefore the interpersonal differences in lifetime incomes) – but this bias is only limited. The appendix documents also that the estimation method adopted in this paper generates plausible life-cycle profiles of income and socioeconomic characteristics by comparing these profiles with cross-section data.

[4] 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 also find that the contribution of social insurance to the redistribution of lifetime incomes 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).

The construction of synthetic life cycles through statistical matching of data for different individuals is not the only method for estimating lifetime incomes. An alternative method – applied by Falkingham et al. (1993), Harding (1993), Nelissen (1998) and Pettersson and Pettersson (2003) among others – is to employ a dynamic microsimulation model to generate a set of representative life cycles, using estimated probabilities to model each individual's transition from one socioeconomic status to another. In this type of model, the evolution of a person's income over the life cycle is simulated as a stochastic process that depends on a number of socioeconomic characteristics. The parameters of the income-generating processes are estimated by regression analysis or fixed by calibration.

Compared to the method of statistical matching adopted here, the use of a microsimulation model has the advantage that it avoids the need to match data for different individuals with various unobservable characteristics. The disadvantage of microsimulation is that the data are artificial and that the quality of the generated data depends on the accuracy of the estimated transition probabilities and income-generating processes. Comparing the two alternative methods for the construction of lifetime incomes is beyond the scope of this paper.

2.2. The redistributive impact of social-insurance transfers

Based on our estimates of lifetime incomes, Table 1 summarizes the redistributive effects of the most important social-insurance transfers in the Danish welfare state, measured by the 'redistribution index'. The calculation of the redistribution index starts from 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, which is a commonly used measure of redistribution, 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 1 normalizes the redistribution index by reporting the *excess* value of the index above the value of the redistribution index for an identical lump-sum transfer to all individuals. The numbers in the first two columns in Table 1 thus indicate how much more redistributive the various transfers are compared to a uniform lump-sum transfer.

It should be stressed that Table 1 measures only the degree to which the various transfer programs systematically redistribute across various income groups. Even if no such redistribution took place, a given transfer could still result in significant redistribution in other dimensions of interpersonal heterogeneity; for example, child benefits obviously result in a systematic redistribution from childless citizens to families with children. Hence, a low value of the redistribution index in Table 1 does not necessarily mean that the transfer program in question fails to achieve its intended distributional goals.

(Table 1 about here)

The first column in Table 1 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 have a smaller impact on income distribution. Disability benefits represent the most important exception to this rule, being much more redistributive in a lifetime context. The reason is that this transfer program offers a relatively high replacement rate; therefore, measured by annual income, the disabled are not among the poorest income groups. However, since disability implies a highly persistent negative income shock, the disabled tend to end up with relatively low lifetime incomes. In a lifetime perspective, therefore, disability benefits are more redistributive. Moreover, a considerable part of housing benefits is granted to recipients of disability benefits. This helps to explain why housing benefits are slightly more redistributive in a lifetime context than in an annual context. Unemployment insurance benefits are also a bit more redistributive in a lifetime perspective.

This is because the incidence of long-term of unemployment tends to be concentrated on unskilled groups whose lifetime incomes are relatively low.

In general, the ranking of the various transfers according to their redistributive impact changes significantly as one moves from an annual to a lifetime measure of income. 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) have a significant impact on the distribution of annual incomes, but exert (almost) the same effect on the distribution of lifetime incomes as an identical lump-sum transfer to all individuals. The reason is that these benefits are granted in a phase of the life cycle when people earn low annual incomes, thereby helping to reduce inequality in annual incomes. However, the individuals who collect these benefits enjoy higher incomes in other seasons of their life course, so these benefits do not contribute much towards narrowing differences in lifetime incomes. The same type of argument holds even more strikingly for grants to students in higher education. Whereas such grants are highly redistributive in an annual context, they generate only small effects on the distribution of 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 at the 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.

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

3.1. 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 help to keep track of the individuals that 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 employs 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 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 labour supply. Rather than dropping its redistributional objectives, the government can achieve these objectives in a more efficient way if it implements individual accounts with a lifetime income guarantee.

In the context of a simple model with two types of households, Box 2 shows how individual saving accounts can improve the equity-efficiency trade-off by using information on lifetime incomes. This additional information essentially allows the government to implement an optimal non-linear lifetime income tax and implement self-insurance for high-income (and middle-income) workers. In particular, individual accounts establish an efficiency-enhancing actuarial link between contributions and benefits for those who end up with a surplus on their account – who currently pay distortionary taxes partly to finance distortionary social benefits to themselves.

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 Box 2). 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, it 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 to (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 considerably less skewed than 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. With many women participating in the labour force and many individuals engaging in long periods of full-time or part-time education, a substantial number of persons move between periods of full-time work to periods of voluntary (sometimes part-time) absence from the labour market to educate themselves, start up a 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 at the beginning of their working lives. 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.

(Start of Box 2 about here)

Box 2. Lifetime redistribution and insurance via individual accounts

To illustrate in the simplest possible manner how individual accounts (IAs) allow the government to redistribute lifetime income in a more efficient manner, consider an economy inhabited by a high-income earner (indicated by superscript h) and a low-income earner (indicated by superscript l). Both individuals live for two periods. In the second period they are retired, and in the first period they are unemployed part of the time and work for the rest of the time. By increasing their job search effort (denoted by a), both persons can reduce their rate of unemployment ($u(a)$) in the first period. Whereas the marginal returns to search efforts (in terms of a lower risk of unemployment) are declining with the level of this effort a, the marginal utility costs are rising in search effort. With C_1 and C_2 denoting consumption in the two periods, and setting the interest rate and the utility discount rate equal to zero to simplify exposition (this assumption is quite innocent), we write the lifetime utility of the high-income earner is

$$U^h = v\left(C_1^h\right) + v\left(C_2^h\right) - f\left(a^h\right), \qquad v' > 0, \quad v'' < 0, \qquad f' > 0, \quad f'' > 0, \tag{2.1}$$

where $f(a)$ is the disutility of search effort. Assuming a perfect capital market for the moment, denoting saving by S, and recalling that the interest rate is zero, we can write consumption in the two periods as

$$C_1^h = W(1-\tau-s)\left[1-u\left(a^h\right)\right] + u\left(a^h\right)B - S^h, \qquad u' < 0, \quad u'' \geq 0, \tag{2.2}$$

$$C_2^h = S^h + y + A^h. \tag{2.3}$$

The variable W is the wage rate of the high-income earner, τ is the labour income tax rate, B is the rate of unemployment benefit, y is a public retirement benefit, s is the rate of mandatory social security contribution to the individual account, and A^h is the balance on the account that is paid out at the time of retirement. If a fraction α of unemployment benefits received is debited to the individual account, the account balance at retirement amounts to

$$A^h = sW\left[1-u\left(a^h\right)\right] - \alpha u\left(a^h\right)B, \qquad 0 \leq \alpha \leq 1, \qquad A^h \geq 0. \tag{2.4}$$

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 y. Using (2.2) through (2.4) and assuming that the constraint $A^h \geq 0$ is not violated, we obtain the consolidated lifetime budget constraint

$$C_1^h + C_2^h = W(1-\tau)\left[1-u\left(a^h\right)\right] + (1-\alpha)\,u(a^h)B + y. \tag{2.5}$$

The social security contribution rate s has dropped out of (2.5), since contributions to the IA are effectively remitted to the consumer when the account balance is paid out. Equation (2.5) also shows that, for a consumer with a surplus on the IA, the account system reduces the effective rate of unemployment benefit by the fraction α.

The high-income earner chooses his search effort a^h so as to maximize lifetime utility (2.1) subject to the budget constraint (2.5). It is easy to show that this behaviour implies

$$\frac{\partial a^h}{\partial \tau} < 0, \qquad \frac{\partial a^h}{\partial B} < 0, \qquad \frac{\partial a^h}{\partial \alpha} > 0, \tag{2.6}$$

given our assumptions of increasing marginal disutility of search ($f'' > 0$) and non-increasing marginal returns to search effort ($u'' \geq 0$). The solution to the high-income earner's problem yields an indirect utility function of the form $V^h(\tau, B, \alpha, y)$ with partial derivatives

$$V_\tau^h = -\lambda^h W\left(1-u^h\right), \qquad V_B^h = \lambda^h(1-\alpha)\,u^h, \qquad V_\alpha^h = -\lambda^h B u^h, \tag{2.7}$$

where λ^h is the high-income earner's marginal utility of (exogenous) income. Finally, we may write the generational account of a high-income earner (i.e., the present value of the net taxes paid over his life cycle) as

$$\begin{aligned} G^h &= (\tau+s)\,W\left[1-u\left(a^h\right)\right] - Bu\left(a^h\right) - y - A^h \\ &= \tau W\left[1-u\left(a^h\right)\right] - (1-\alpha)\,Bu\left(a^h\right) - y. \end{aligned} \tag{2.8}$$

The low-income earner has a utility function analogous to (3.1), but his income is assumed to be so low relative to the unemployment benefits he received that he ends up with a deficit on his IA. Hence, his account is set to zero at retirement, so that the mandatory contributions paid into the account of the low-income earner are not remitted to him. Instead, these contributions will work like the ordinary labour income tax, as seen from the following lifetime budget constraint of the low-income earner:

$$C_1^l + C_2^l = w(1-\tau-s)\left[1-u\left(a^l\right)\right] + u\left(a^l\right)B + y, \tag{2.9}$$

where $w < W$ is the low-income earner's wage rate. With optimal search effort, the low-income earner's indirect utility function can be shown to take the form $V^l(\tau, s, B, y)$ with derivatives

$$V^l_\tau = V^l_s = -\lambda^l w \left(1 - u^l\right), \qquad V^l_B = \lambda^l u^l, \tag{2.10}$$

and his generational account is

$$G^l = (\tau + s)\, w \left[1 - u\left(a^l\right)\right] - Bu\left(a^l\right) - y. \tag{2.11}$$

We now show that, starting from a situation without individual accounts where $s = \alpha = 0$, the introduction of IAs (i.e., the introduction of positive values of s and α) for a particular cohort enables the government to generate a Pareto improvement (the older cohorts are not affected, since the original tax-transfer system remains in place for them). To see this, note from (3.7) and (3.10) that the policy reform

$$d\tau = -ds < 0, \qquad d\alpha = \frac{W\left(1 - u^h\right)}{Bu^h} \cdot ds > 0, \tag{2.12}$$

will ensure that $dV^l = dV^h = 0$; that is, it will keep the (ex-ante) lifetime utilities of both individuals constant. If at the same time the policy reform (2.12) increases net public revenue (by increasing G^h and/or G^l), it must be possible for the government to make everybody better off – say, by using the extra revenue to raise the retirement benefit y. It is immediately clear from (2.9) and (2.11) that the policy $d\tau = -ds$ affects neither the behaviour nor the generational account of the low-income earner, since it leaves his total effective tax rate $\tau + s$ unchanged. Also the generational accounts of older generations are not affected, since these cohorts remain under the old system. For the high-income earner, however, we find from (2.8) that

$$dG^h = (\tau W + B) \left[\overbrace{\left(\frac{\partial u^h}{\partial \tau}\right)}^{>0} ds - \overbrace{\left(\frac{\partial u^h}{\partial \alpha}\right)}^{<0} d\alpha \right] > 0, \tag{3.13}$$

where we have used (2.6) and $\partial u^h / \partial x = u'(a^h) \cdot \left(\partial a^h / \partial x\right)$. The introduction of IAs thus improves the public budget through two channels. First, the closer actuarial link between contributions and benefits reduces the effective labour-income tax rate facing the high-income earner, thereby raising search effort. Second, the individual accounts reduce the effective rate of unemployment benefit that the high-income earner enjoys. This induces him to reduce the risk of unemployment by searching more intensely for a job. Less moral hazard in unemployment insurance thus results in a lower unemployment rate.

(End of Box 2 here)

3.2. Moral hazard and optimal lifetime insurance

Another reason why saving schemes may enhance efficiency involves 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.[5]

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, the 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 labour market to enjoy leisure, educate themselves, set up a business, or care for children or frail relatives. In such a transitional labour 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 those 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

[5] Whereas the premia of such insurances do not distort the labour market, the benefits harm labour-market incentives if people can affect the probability that they are eligible for these benefits by changing their labour-market behaviour.

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 agents and to regulate their behaviour, 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 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 each of half a year reduce lifetime incomes of an individual with a full-time working career of thirty years by only about 3%.

The accounts in fact combine a number of risks that occur during different periods of an individual's life in a single insurance contract with a deductible that is conditioned on the aggregate loss during the life course. Drawing on Arrow (1971), Gollier and Schlesinger (1995) show that an umbrella insurance contract with a deductible on the aggregate loss dominates separate insurance contracts with separate deductibles for each type of loss if these individual losses are imperfectly correlated. Intuitively, insurance protection is most valuable for the states in which several losses occur simultaneously. An umbrella insurance policy that adjusts the deductible on each separate loss to the outcome of the other risks in the form of a straight deductible based on the aggregate loss provides the best protection against large aggregate losses for a given insurance budget. Hence, an individual account that insures aggregate lifetime risk with a single deductible conditioned on aggregate lifetime losses dominates separate insurance contracts with their own deductibles for risks faced in different periods of an individual's lifetime. For a given level of insurance cover, separate policies for risk experienced in each period underindemnify risk-averse individuals for high levels of aggregate losses during the life course and overindemnify them for losses that are experienced only in a single period. Compared to separate insurance policies, the umbrella insurance contract provides better protection in the worst-case scenario of a succession of adverse shocks during the life course in exchange for less protection in other cases.[6]

On the basis of a simple model, Box 2 indicates how individual accounts can combat moral hazard in unemployment insurance. This box shows that the individual 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 where people remain protected by the lifetime income guarantee. This is in fact an application of Gollier and Schlesinger (1995) to the insurance of human capital. The government can insure individuals more efficiently against unemployment and lack of skills by basing the deductible in lifetime income unsurance on aggregate income losses during the life course instead of separate deductibles for unemployment in a single period. The model in Box 2 shows that an umbrella insurance implemented by the individual accounts is more efficient than separate insurances for each separate contingency. In particular, if the government keeps expected utilities constant, it can reap the gains from more efficient insurance in terms of additional budgetary resources as a result of higher employment and lower unemployment. The low-skilled who suffer from unemployment and low labour incomes when employed retain the original unemployment protection. The high-skilled, in contrast, are confronted with a higher deductible when becoming unemployed. If the government exactly compensates these high-skilled individuals for the larger exposure to unemployment risk by allowing them to keep part of their contribution to the unemployment insurance system as higher retirement benefits, then public revenues will increase under two conditions.[7] First, search effort should respond to

[6] We are indebted to Christian Gollier for pointing out this link to the literature on optimal insurance.

[7] A third condition is that agents are not completely myopic and thus value an increase in future

incentives so that moral hazard is present in unemployment insurance. Second, not only the lowest skill categories but also other agents should suffer the risk of unemployment so that the risk of being born low skilled is not perfectly correlated with the risk of becoming unemployed.

More generally, 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 the 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, labour-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 labour 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

retirement benefits as a result of more current contributions into their accounts.

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 each industrial sector determining its own tailor-made mix 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.

3.3. 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 in effect 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. If it replaces regular public transfers by individual accounts, the government thus does not abandon its role in alleviating liquidity constraints. Indeed, even with imperfect capital markets, individual saving accounts allow individuals to smooth consumption over their life cycle. By helping 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 unemployment, Box 3 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 where 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 the debiting of benefits to individual unemployment accounts would allow workers to concentrate the cut in consumption in periods when they are employed and have a lower marginal utility of consumption.

(Start of Box 3 about here)

Box 3. Efficient liquidity insurance via individual accounts

To highlight how the introduction of IAs allows the government to offer liquidity insurance to the unemployed in a more efficient manner, we focus on the high-income earner introduced in Box 3. However, now we assume that he is liquidity-constrained in period 1, thus consuming all of his disposable income in that period, and that he is fully employed in period 2. Using the notation of Box 2, but abstracting from the public retirement benefit, we can write the high-income earner's consumption in the two periods of life as

$$C_1^h = W(1-\tau-s)\left[1-u\left(a^h\right)\right]+u\left(a^h\right)B, \tag{3.1}$$

$$C_2^h = W(1-\tau-s)+A^h = W(1-\tau)+sW\left[1-u\left(a^h\right)\right]-\alpha u\left(a^h\right)B, \tag{3.2}$$

where we have used the specification of the account balance A^h stated in equation (2.4) of Box 2. The high-income earner maximizes his utility function (2.1) in Box 2 with respect to his search effort a^h, subject to (3.1) and (3.2). The solution to this problem yields an indirect utility function $V^h(\tau, s, B, \alpha)$ with partial derivatives

$$V_\tau^h = -W\left[\lambda_2^h+\lambda_1^h\left(1-u^h\right)\right], \qquad V_B^h = \lambda_1^h u^h, \qquad V_\alpha^h = -\lambda_2^h u^h B, \tag{3.3}$$

where λ_1^h and λ_2^h are the marginal utilities of income in the two periods, and where the binding liquidity constraint implies that $\lambda_1^h > \lambda_2^h$ (the expression for V_B^h assumes that the initial value of α is zero).

Under a conventional tax-transfer system without IAs ($s = \alpha = 0$), the government can alleviate the liquidity constraint through a combined rise in the rate of unemployment benefit B and the labour-income tax rate τ. Under a system with IAs, an alternative way of relaxing the liquidity constraint would be to combine the rise in B with a rise in α, thus debiting (part of) the higher unemployment benefit to the IA, thereby forcing the consumer to self-finance the higher benefit at a later stage in life when the marginal utility of income is lower.

We will now show that the latter procedure involving the use of unemployment accounts provides liquidity insurance to the unemployed at a lower efficiency cost than the conventional tax-financed benefit system. In demonstrating this, we shall use the high-income earner's generational account, which is now given by

$$\begin{aligned} G^h &= \overbrace{(\tau + s) W \left[1 - u\left(a^h\right)\right] - u\left(a^h\right) B}^{\text{net taxes paid during period 1}} + \overbrace{(\tau + s) W - A^h}^{\text{net taxes paid during period 2}} \\ &= \tau W \left[2 - u\left(a^h\right)\right] - u\left(a^h\right) B (1 - \alpha), \end{aligned} \tag{3.4}$$

where we have maintained the simplifying assumption from Box 3 that the discount rate is zero.

Suppose initially that there is no IA system ($s = \alpha = 0$), and consider a conventional tax-transfer policy involving a combined rise in B and τ, which is calibrated to keep the lifetime utility of the high-income earner constant, that is:

$$dV^h = V_\tau^h \cdot d\tau + V_B^h \cdot dB = 0 \quad \Longrightarrow \quad d\tau = \left(\frac{\lambda_1^h u^h}{W \left[\lambda_2^h + \lambda_1^h (1 - u^h)\right]} \right) \cdot dB. \tag{3.5}$$

With constant consumer utility, this policy will imply a gain (loss) in economic efficiency if it increases (reduces) public net revenue. From (3.4), we find by using (3.5) that the combined rise in B and τ has the following revenue effect:[8]

$$\frac{dG^h}{dB} = \overbrace{\left(\frac{u^h \left(\lambda_1^h - \lambda_2^h\right)}{\left[\lambda_2^h + \lambda_1^h (1 - u^h)\right]} \right)}^{\text{liquidity insurance effect } (>0)} \overbrace{-u' \cdot (B + \tau W) \left[\frac{\partial a^h}{\partial B} + \left(\frac{\partial a^h}{\partial \tau} \right) \left(\frac{d\tau}{dB} \right) \right]}^{\text{incentive (employment) effect } (<0)}. \tag{3.6}$$

[8] ds can be set so that the utility of the low-income earner changes in exactly the same way as in the case in which the government finances the higher unemployment benefits through higher taxes. This requires that the increase in the contribution rate with individual accounts ds is equal to the increase in the marginal tax rate $d\tau$ with regularly financed unemployment benefits.

The first positive term on the right-hand side of (3.6) measures the welfare gain from improved liquidity insurance, reflecting the fact that a combined rise in B and τ tends to shift disposable income from the second period towards the first period of life where the marginal utility of income is higher. The second large term on the right-hand side of (3.6) captures the efficiency loss from weaker incentives for job search. Using the consumer's first-order condition, one can show that $\partial a^h/\partial B < 0$, i.e., the higher unemployment benefit reduces job search intensity. In principle, the effect of a higher tax rate on first-period job search is ambiguous because of offsetting income and substitution effects. However, on the standard assumption that a higher tax rate has a non-positive effect on search effort, we have $\partial a^h/\partial \tau \leq 0$, which is sufficient to ensure a negative employment effect of the policy (3.5). Hence, we cannot say a priori whether that policy has a positive net impact on public revenue and economic welfare.

Consider again an increase in unemployment benefits, but suppose now that instead of financing it through a higher labour income tax, the government introduces individual unemployment accounts, combining the rise in B with a rise in α (from zero to some positive number) such that

$$dV^h = V_\alpha^h \cdot d\alpha + V_B^h \cdot dB = 0 \quad \Longrightarrow \quad d\alpha = \left(\frac{\lambda_1^h}{\lambda_2^h B} \right) \cdot dB. \tag{3.7}$$

Since this policy also keeps consumer utility constant, it will be more efficient than the conventional tax-transfer policy (3.5) if it has a more positive (or less negative) impact on the public budget. From (3.4) and (3.7) we find

$$\frac{dG^h}{dB} = \overbrace{\left(\frac{u^h \left(\lambda_1^h - \lambda_2^h \right)}{\lambda_2^h} \right)}^{\text{liquidity insurance effect } (>0)} \overbrace{-u' \cdot (B + \tau W) \left[\frac{\partial a^h}{\partial B} + \left(\frac{\partial a^h}{\partial \alpha} \right) \left(\frac{d\alpha}{dB} \right) \right]}^{\text{incentive (employment) effect } (\gtreqless 0)}. \tag{3.8}$$

Again, we have a positive liquidity insurance effect, but compared to the conventional policy generating the revenue effect in (3.6), we see from (3.8) that the liquidity insurance effect is stronger under the IA system. The reason is that under the conventional tax-transfer policy (3.5), the positive liquidity effect of higher unemployment benefits in the liquidity-constrained period 1 is partly offset by the rise in the labour income tax. Furthermore, one can use the consumer's first-order condition to show that $\partial a^h/\partial \alpha > 0$; thus, from (3.8) we see that the negative impact of the higher unemployment benefit on job search is (partly) offset by the fact that the consumer (partly) self-insures against unemployment via the IA. By contrast, when insurance takes place through the tax system, we observe from (3.6) that the higher tax rate will tend to have a negative employment effect. For these reasons, an IA system offers liquidity insurance in a more efficient manner than does a conventional tax-transfer system.

(End of Box 3 here)

3.4. 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. If the government replaces regular public benefits by individual accounts with a lifetime income guarantee, the government thus not only maintains liquidity insurance and lifetime redistribution but also continues to assist individuals in smoothing consumption over the life cycle. To the extent that the government does not fully succeed in helping individuals who lack self control to implement consumption smoothing, annual disposable income becomes more important compared to 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.

Another cost of mandatory saving is that it harms labour-supply incentives if people are myopic or suffer from liquidity constraints. The positive incentive effects from introducing individual accounts in Box 2 rely on individuals being forward looking and valuing the future benefits (in terms of higher future retirement benefits) from searching harder today. The accounts would have exactly the same disincentive effects as tax-financed unemployment benefits if agents were completely myopic.

3.5. Administration

The individual accounts can be administered in various ways. One option is to give individuals a choice about which private financial institution will administer their accounts. On the one hand, the competitive pressure may encourage cost effectiveness and innovation in the administration of the accounts. On the other hand, competition may give rise to additional marketing, selling and transaction costs. To contain these costs, the government may want to administer the accounts itself. In that case, individuals obviously have no choice in who administers their accounts.

3.6. Comparison with alternatives

Table 2 compares mandatory individual accounts that are administered by the public sector and involve a lifetime-income guarantee to three alternatives: voluntary saving, a Bismarckian insurance system involving a clear actuarial link at the individual level between premia paid and the value of the insurance provided, and a redistributive transfer system in the tradition of Beveridge with flat benefits but progressive taxation.

Compared to voluntary saving, mandatory individual accounts redistribute, offer liquidity insurance, protect individuals lacking self control, and limit transaction and marketing costs. Just as voluntary saving, they combat moral hazard and limit labour-supply distortions for forward-looking middle- and high-income earners. However, the accounts fail to achieve this for the lifetime poor. This is in contrast to voluntary saving.

(Table 2 about here)

The accounts share with Bismarckian insurance the benefits of liquidity insurance, limited transaction costs, and protection of individuals who lack self control. They differ from this type of insurance in two important respects. First, the accounts redistribute between the lifetime poor and the lifetime rich through a lifetime income guarantee by bailing out those individuals who end up with a negative balance at retirement. The price of this redistribution is that they do not provide an actuarial link for the lifetime poor, thereby harming labour-market incentives for this group. The second difference between Bismarckian insurance and the accounts is that the accounts fight moral hazard because insurance benefits are taken out of the individual accounts. The other sign of the coin is that, compared to Bismarckian insurance, the accounts provide less good insurance for those people who end up with a positive account balance.

3.7. 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. 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); and 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 the income distribution, the public finances, and economic efficiency. As indicated by the discussion in the previous section and by the analysis in Sørensen (2003) and Bovenberg and Sørensen (2004), a system of IAs with a lifetime-income guarantee could potentially generate an ex-ante Pareto improvement because it represents an additional policy instrument that helps to alleviate the equity-efficiency trade-off. Section 4.1 presents a simple method of estimating whether a specific IA system would in fact be Pareto-improving in an ex-post sense (i.e. not only in an expected utility sense before people know which shocks will hit them during their lives, but also in terms of realized utility in each contingency). To illustrate the policy choices involved in the design of an IA system, section 4.2 then considers a concrete proposal for a system of individual accounts in Denmark. Section 4.3 estimates the effects of the proposed IA system on the distribution of lifetime incomes, and section 4.4 uses the method developed in section 4.1 to explore its likely effects on the labour market, the public budget and economic efficiency.

4.1. Estimating the revenue and welfare effects of individual accounts

Under the proposed system of individual accounts, a part of the individual taxpayer's current tax bill is converted into a mandatory social security contribution that is credited to his IA, while transfers from the programs included in the system are debited to his IA. The individual account balance A_t accumulates at the (after-tax) market interest rate r_t, so if a fraction s of labour income Y_t is contributed to the account during period t, the account balance at the start of time $t+1$ is

$$A_{t+1} = (1 + r_t) A_t + sY_t - \alpha B_t, \quad (1)$$

where B_t is a column vector representing the amounts of the various public transfers received by the account holder during period t, and α is a row vector indicating the fractions of the various transfer payments that are debited to the IA.

To receive the transfers included in the IA system, the proposal considered here requires that the account holder must satisfy the same eligibility criteria as under the present transfer programs. If the account holder has a negative account balance when he/she reaches the official retirement age, the balance is simply set to zero, and the holder receives the normal public pension in accordance with the current rules. If the account balance is positive, it is converted into an annuity, which is added to the owner's ordinary public pension. Thus, the IA system has a lifetime-income guarantee ensuring that no person can be worse off in any realized contingency than under the existing tax-transfer system, while individuals with positive account balances will be better off.[9] If the introduction of such a system can be shown to be self-financing (or even to improve the public budget), then it implies a Pareto improvement in terms of realized utilities if non-fiscal external effects are absent.[10]

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

[9] This procedure is different from that used in Box 2, where we set the parameters of the account system in such a way that all agents are indifferent between the account system and the initial tax-transfer system in terms of expected utilities. In the present section, we allow individuals with a surplus on their accounts to be better off than under the initial tax-transfer system. Another difference is that Box 2 explores the effects on ex-ante utility (i.e. before people know what shocks will hit them) while in the present section we explore the impact on ex-post utilities after shocks have in fact materialized.

[10] The complications arising when these latter external effects are present will be discussed below.

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, 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.

Box 4 offers a formula that can be used to estimate the budgetary effect of an IA system of the type just described under the assumption that all cohorts born later than some specific year would be subject to the IA system, whereas all earlier cohorts would be subject the current tax-transfer system. The system would thus be phased in gradually without any intergenerational redistribution effects. The formula derives the impact of the reform on net public revenue from the cohorts that fall under the new system.

Box 4 focuses on changes in public expenditure and in revenues from (direct and indirect) taxes on labour. For simplicity, it thus abstracts from any changes in the revenue from capital income taxes that might follow from changes in saving behaviour. However, aggregate saving is unlikely to change much since intergenerational redistribution effects are absent in our account system. Indeed, the individual accounts are structured so that the impacts on the inter- and intragenerational distribution are minimized. Similarly, the proposal does not yield any first-order implications for the impact of capital-market imperfections such as liquidity constraints. Our saving accounts allow individuals to enjoy the same insurance benefits as under the current tax-transfer system – even if their account balance is negative. In this respect, the accounts provide the same liquidity insurance as current public benefits. Furthermore, by forcing working individuals to put part of their labour income in the mandatory saving account (instead of in the public accounts), the government helps to ensure that the impact of the accounts on private consumption and saving behaviour is minimized.

The formula in Box 4 rests on the simplifying assumption that the interest rate paid on account balances roughly equals the rate of income growth so that the growth-adjusted interest rate is zero. Note that the interest rate applied to the IA balances should be the *after-tax* rate of interest to ensure that saving via the IAs is treated in the same manner as ordinary saving via the capital market. As illustrated in Figure 2, the average after-tax interest rate on long-term government bonds has in fact been fairly close to the growth rate of GDP in Denmark in recent years.

(Figure 2 about here)

If the net interest rate would be larger (smaller) than the growth rate, the various terms in formula (4.6) in Box 4 would be scaled up (down) by an additional factor, depending on when the various revenue effects identified in the formula occur. This additional factor is the same for all terms in the formula if the average duration between contributions to the accounts and the time of retirement coincides with the corresponding average duration for withdrawals. This condition is not a bad approximation for social insurance benefits collected during the working life.[11] If this condition holds, all terms in the formula would be scaled up (if the net interest rate exceeds the growth rate) or down (if the after-tax interest rate is lower than the growth rate) by the same factor. Hence, the overall revenue effect would have the same sign as with a zero growth-adjusted interest rate. The absolute size of the effect, however, would be larger or smaller with a percentage equal to the growth-adjusted interest rate times the average duration.

The condition that the durations of contributions and withdrawals are the same would also be relevant under an alternative transition strategy under which all citizens between a certain age (say, 18) and the official retirement age would start participating immediately. Under this transition strategy, people reaching the retirement age would have accumulated only relatively small balances during the first few years. As time elapses and new retirees would have participated in the IA system for a longer time, balances would grow. Larger positive balances, which would be channeled back to the retirees, would be a drain on the public finances, but at the same time the labour-supply effects discussed below would exert a growing positive influence on net public revenue derived from the retiring generations. Under the condition that the average duration between contributions and the time of retirement coincides with the corresponding average duration for withdrawals, the various terms of the formula in Box 4 for revenues collected

[11] In the case of retirement benefits, in contrast, contributions are collected earlier than withdrawals are paid so that the duration of contributions exceeds that of withdrawals. In that case, the account balances and thus the negative impact of the introduction of the accounts on the public finances would increase if the interest rate would exceed the growth rate. The static revenue losses thus would increase compared to the other terms in formula 4.6 in Box 4.

from generations would all grow as cohorts would get younger and would thus be subject to the account system for a longer time. Hence, the overall revenue effect for the transitional generations would have the same sign but would increase in absolute value until the generations that have been subject to the account system during their entire working life would start to retire.

(Start of Box 4 about here)

Box 4. Estimating the effects of individual accounts on the public budget

In deriving the budgetary impact of the IA system described in section 4.1, we focus on the change in the revenue from labour income taxes and indirect taxes and the change in government expenditure on transfers. Under the assumption that the risk-free real interest rate equals the economy's real (trend) growth rate, we may measure all magnitudes in current income levels and add up net public revenues in different time periods to arrive at net revenue measured in growth-adjusted present-value terms. Since persons who end up with a deficit on their IA pay the same taxes and receive the same transfers as under the current system, we may focus on those individuals who manage to accumulate a surplus on their IA at the date of retirement. With a zero growth-adjusted real interest rate, and with the total time available up until the official retirement age normalized at unity, the balance (A) on the IA at that date will be

$$A = sewh - \alpha_e b_e(1-e) - \alpha_m b_m m, \qquad 0 \leq \alpha_e; \quad \alpha_m \leq 1, \tag{4.1}$$

where s is the rate of mandatory contribution to the IA; w is the wage rate of the representative wage earner with an IA surplus, and e and h are his average labour-force participation rate and working hours, respectively (so that ewh is his total labour income); b_e is his average after-tax public transfer received in periods of non-employment, and α_e is the fraction of out-of-work benefits during working age that is debited to the IAs for this person. The variable b_m is another public benefit for this representative wage earner that is not directly tied to non-employment (e.g. child benefits; education benefits) but to another variable m (e.g. the number of children, time spent on education). The parameter α_m is the fraction of these benefits that is debited to the IAs (again, for those with positive balances), and we allow for the possibility that the variable m (e.g. education) may affect employment (i.e. $de/dm \neq 0$). With T denoting the average tax payment over the active career of a person with an IA surplus, the growth-adjusted present value of the net public revenue collected from the representative person with an IA surplus is

$$\begin{aligned} R &= eT + sewh - (1-e)\, b_e - b_m m - y - A \\ &= eT - (1-\alpha_e)\,(1-e)\, b_e - (1-\alpha_m)\, b_m m - y, \end{aligned} \tag{4.2}$$

where y is the ordinary retirement benefit granted to people above the official retirement age (assumed not to be included in the IA system). The second line in (4.2) shows that the contribution rate s yields no revenue effects, because all contributions are effectively remitted to individuals with a positive IA balance. It also highlights that under an IA system, people with an IA surplus self-finance a fraction α of the benefits they receive during their working lives. Finally, note that all variables in (4.2) are measured after indirect taxes, so the revenue effects of indirect consumption taxes are implicitly included (see the specification of effective tax rates in section 4.3).

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

$$T = \tau wh - I. \tag{4.3}$$

Here, τ is the marginal 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 income.

The introduction of IAs means that part of the labour income tax is replaced by a mandatory IA contribution, and that part of the transfers received is debited to the IA. In formal terms, such a reform thus implies a cut in τ combined with a rise in the variables s, α_e and α_m. We wish to estimate the revenue effect of introducing a system of IAs, starting from an initial situation without such a system where $A = s = \alpha_e = \alpha_m = 0$. Using this initial condition, we show in an appendix available at www.econ.ku.dk/pbs/ that the revenue effect of introducing IAs amounts to

$$\frac{dR}{ewh} = \overbrace{d\tau + c_e\left(\frac{1-e}{e}\right) d\alpha_e + c_m\left(\frac{m}{e}\right) d\alpha_m}^{\text{static effect}}$$

$$-\underbrace{\varepsilon\left(\frac{\phi_h\tau}{1-\phi_h\tau}\right)\bar{\phi}_h d\tau}_{\text{hours-of-work effect}}\overbrace{}^{\text{dynamic effect}}-\underbrace{\eta\left(\frac{\phi_e t+c_e}{1-\phi_e t-c_e}\right)\bar{\phi}_e d\tau}_{\text{participation effect}}$$

$$\overbrace{+\underbrace{(\phi_e t+c_e)\left\{\left(\frac{c_e\eta}{1-\phi_e t-c_e}\right)d\alpha_e+\varphi\chi d\alpha_m\right\}}_{\text{co-insurance participation effect}}+\underbrace{\chi c_m\left(\frac{m}{e}\right)d\alpha_m}_{\text{co-insurance moral-hazard effect}}}^{\text{dynamic effect}} \tag{4.4}$$

where $c_e \equiv b_e/wh$ and $c_m \equiv b_m/wh$ are, respectively, replacement rates and normalized benefit rates in the transfer programs included in the IA system for those who end up with an IA surplus, $t \equiv T/wh$ denotes the average labour income tax rate, ε is an hours-of-work elasticity indicating how hours worked respond to a change in the marginal after-tax wage rate, η is a participation elasticity reflecting how labour-force participation reacts to a change in the difference between net income from employment and net income from non-employment, $\varphi \equiv -\frac{de}{e}/\frac{dm}{m}$ is the (negative) elasticity of employment with respect to our variable m (e.g., the elasticity of employment with respect to education), and χ is a 'benefit dependency elasticity' reflecting how the eligibility criterion m responds to a change in the net benefit rate $b_m(1-\alpha_m)$. The parameters ϕ_h and ϕ_e measure the extent to which the (marginal) labour income tax τ is a 'genuine' tax rather than an insurance premium entitling the taxpayer to additional social insurance benefits. Under a pure Bismarckian social insurance system with an actuarially fair link between social security taxes paid and benefits received, we would have $\phi_h = \phi_e = 0$. At the opposite end of the spectrum, under a pure Beveridgean social insurance system with no link between taxes and benefits at the individual taxpayer level, we have $\phi_h = \phi_e = 1$. The parameters $\bar{\phi}_h$ and $\bar{\phi}_e$ in (5.4) are defined as $\bar{\phi}_h \equiv d(\phi_h\tau)/d\tau$ and $\bar{\phi}_e \equiv d(\phi_e\tau)/d\tau$, thus measuring the degree to which the marginal effective tax rates $\phi_h\tau$ and $\phi_e\tau$ vary with the statutory marginal labour-income tax rate τ. These parameters may depend on the programs that are included in the individual accounts.

The 'static' effect indicated in (4.4) measures the immediate impact on net public revenue of introducing IAs, abstracting from any change in individual behaviour. However, the introduction of IAs and the associated cut in the labour income tax will affect labour-force participation, hours worked, and the take-up of social benefits, and these behavioural responses will generate the 'dynamic' budgetary effects indicated in (4.4). The 'co-insurance participation effect' reflects that, via the IA system, consumers partly self-finance the social benefits they receive during periods of non-employment. This induces them to reduce their periods of non-employment, thereby strengthening the public finances. Similarly, the (partial) self-financing of benefits via the IAs may induce some individuals to take up to a smaller extent certain benefits that are not directly related to non-employment, as reflected in the 'co-insurance moral hazard effect' in equation (4.4).

Following the proposal presented in section 4.2, suppose that the rate of social security contribution s is set so as to ensure that the aggregate contributions to the IAs are equal to the aggregate payments of transfers from the accounts, given the labour income tax base and the expenditure on transfers prevailing before the reform. Recalling that $ds = -d\tau = s$ and $d\alpha_i = \alpha_i$ for $i = e, m$ (since s and α are initially zero), we then have

$$s \cdot e^* w^* h^* = \alpha_e^* \cdot b_e^* (1 - e^*) + \alpha_m^* \cdot b_m^* m^* \Longrightarrow$$

$$-d\tau = c_e^* \left(\frac{1 - e^*}{e^*}\right) \alpha_e^* + c_m^* \left(\frac{m^*}{e^*}\right) \alpha_m^*, \qquad c_e^* \equiv \frac{b_e^*}{w^* h^*}, \quad c_m^* = \frac{b_m^*}{w^* h^*}, \tag{4.5}$$

where the asterisks indicate averages across the entire working population (including those who end up with negative IA balances); c_e^* therefore represents the average replacement rate across the entire labour force. Using (4.5) to eliminate $d\tau$ from (4.4), we obtain

$$\frac{dR}{ewh} = \overbrace{c_e \left(\frac{1-e}{e}\right) \alpha_e - c_e^* \left(\frac{1-e^*}{e^*}\right) \alpha_e^* + c_m \left(\frac{m}{e}\right) \alpha_m - c_m^* \left(\frac{m^*}{e^*}\right) \alpha_m^*}^{\text{static effect}}$$

$$+ \overbrace{\left[\underbrace{\varepsilon \left(\frac{\phi_h \tau}{1 - \phi_h \tau}\right) \bar{\phi}_h}_{\text{hours-of-work effect}} + \underbrace{\left(\frac{\eta(\phi_e t + c_e)}{1 - \phi_e t - c_e}\right) \bar{\phi}_e}_{\text{participation effect}}\right] \underbrace{\left[c_e^* \left(\frac{1-e^*}{e^*}\right) \alpha_e^* + c_m^* \left(\frac{m^*}{e^*}\right) \alpha_m^*\right]}_{\text{cut in labour income tax rate}}}^{\text{dynamic effect from lower tax rate}}$$

$$+ \overbrace{\underbrace{(\phi_e t + c_e) \left\{\left(\frac{c_e \eta}{1 - \phi_e t - c_e}\right) \alpha_e + \varphi \chi \alpha_m\right\}}_{\text{co-insurance participation effect}} + \underbrace{\chi c_m \left(\frac{m}{e}\right) \alpha_m}_{\text{co-insurance moral hazard effect}}}^{\text{dynamic effect from lower effective benefits}} \tag{4.6}$$

(End of Box 4 here)

Box 4 shows that the impact of IAs on the public budget can be decomposed into a 'static' effect reflecting the change in net public revenue that would occur if behaviour were unchanged, and a 'dynamic' effect capturing the budgetary impact of behavioural changes. The static effect is simply the sum total of the positive IA balances that the government would have to transfer to the account holders if they would not change their behaviour. Since these resources were previously part of general government revenue, they measure the deterioration of the public budget in the hypothetical situation where no taxpayer responds to the change in incentives brought about by the IA system. In the absence of behavioral changes, the individual accounts would thus not be Pareto improving in an ex-post sense. In particular, individuals without positive account balances would lose as the government would have to raise taxes to make up for its revenue losses. The result that some individuals would be worse off is different from the result derived in Box 2 that in the absence of behavioral effects the introduction of the accounts would leave all households equally well off in an ex-ante sense. Indeed, the conditions for an ex-post Pareto improvement (where no agent is worse off in any contingencies) are stronger than those for an ex-ante Pareto improvement (where no agent is worse off in terms of expected utility).

Formula (4.6) in Box 4 reveals that the static revenue loss is larger, the more heterogeneous is the population. In particular, the larger is the employment rate and the smaller are the average replacement rate and the call on non-employment benefits (m) for individuals with positive IA balances relative to the population as a whole, the larger will be the positive IA balance accruing to an average person within the group with positive balances, and so the larger will be the revenue loss that occurs in the absence of behavioural responses. Intuitively, with a more heterogeneous population, a conventional tax-transfer system tends to imply more interpersonal redistribution across the life cycle, so the larger is the static budgetary cost of reducing net revenue collection from those individuals who are net taxpayers over the life cycle under the current system. One can also interpret the static revenue loss as the distributional loss from the introduction of individual accounts. This distributional loss is translated into a revenue loss, as the government compensates those who end up losing from the accounts. The distributional loss clearly depends on the type of benefit that is included in the IA system. Some benefits imply larger distributional losses than others.

The dynamic effect, which consists of the product of behavioural effects and tax wedges (or first-order distortions), can be decomposed into two separate effects. The first is the impact of a closer actuarial link between contributions and benefits (see the terms on the second lines of (4.4) and (4.6)). This produces lower effective marginal and average tax rates, resulting in more hours worked and more labour-force participation. This effect depends crucially on the extent to which the lower tax rate τ produces in fact a lower effective marginal and average tax rate on labour, as reflected in the parameters $\bar{\phi}_h$ and $\bar{\phi}_e$ in (4.6). These parameters may depend on the type of contribution (and benefit) included in the accounts. To illustrate, taxes financing child benefits or educational benefits typically do not generate additional benefit rights if one works more hours or participates more. However, in the case of social security contributions for unemployment or sickness benefits (especially in Bismarckian systems), one may obtain additional benefit rights if one moves from being a part-time worker to being a full-time worker. Hence, there may already be some actuarial link here, implying that $\bar{\phi}_h < 1$. Similarly, by entering the labour force, one may become entitled to unemployment, sickness or parental leave benefits. This entitlement effect causes the parameter $\bar{\phi}_e$ to fall below unity.

In addition to the behavioural effects (which are in part determined by the parameters $\bar{\phi}_h$ and $\bar{\phi}_e$ affecting the changes in labour-supply incentives), tax wedges (i.e. $\phi_h \tau$ for hours worked and $(\phi_e t + c_e)$ for labour-force participation; see the second lines of formulas (4.4) and (4.6)) play a role in determining the dynamic revenue effects. The tax wedge on labour-force participation is determined by the average replacement rate for all out-of-work benefits c_e, which may include quite different benefits than those included in the accounts. This average replacement rate thus does not depend on the particular benefits that are included in the accounts, but depends on the overall out-of-work benefits in the economy. The same holds true for the effective tax rate $\phi_e t$, which depends on the nature of all taxes and contributions in the economy rather than the specific contributions that are included in the accounts.

The second component of the dynamic effect is the revenue impact of less benefit dependency on account of more self insurance, as individuals finance their own benefits out of their individual accounts (see the terms on the third lines of (4.4) and (4.6)). This 'co-insurance effect' (or 'benefit dependency effect') is present even in Bismarckian social security systems with actuarially fair contributions (i.e. $\bar{\phi}_h = \bar{\phi}_e = 0$). If the benefit is triggered by non-employment (and thus insures against periods of inactivity), labour-market inactivity will drop. How much this 'co-insurance participation effect' then benefits the budget depends on the initial overall tax wedge on employment $(\phi_e t + c_e)$ and the sensitivity of inactivity with respect to a higher value of α_e, as indicated by the participation elasticity η in formula (4.6).[12] The overall tax wedge depends on both an effective average tax wedge $\phi_e t$ and a replacement rate c_e.

If the benefit is not employment-related (i.e. child benefits), the co-insurance effect does not (directly) affect employment. However, co-insurance may combat moral hazard in other domains. In the case of child benefits, for example, fertility may decline if parents have to pay the child benefits out of their own accounts. Also this 'co-insurance moral-hazard effect' depends on the product of a behavioural elasticity χ and a 'tax wedge' $c_m \left(\frac{m}{e}\right)$ indicating to what extent one less child saves the government money at the margin.

Some benefits may generate an indirect effect on employment. Education benefits may be an example of this. If people have to pay their education benefits out of their own accounts, they may take less education (the budgetary implications of less education are captured by the term $\chi c_m \left(\frac{m}{e}\right) d\alpha_m$ in formula (4.6)). This, however, may yield indirect implications for employment e. For example, if people spend less time in training, they may move earlier to employment, thereby raising employment (implying that $\varphi > 0$).[13] However, less education may also make it less attractive to work longer (implying $\varphi < 0$).The revenue implications of these cross effects on participation depend on the overall tax wedge on participation $(\phi_e t + c_e)$.

The decomposition of the budgetary impact in formula (4.6) into a static and a

[12] We have included only labour-supply effects of lower effective replacement rates in our elasticity η. However, a lower effective replacement rate on account of a higher value for α_e is also likely to reduce wage pressure, thereby boosting labour demand and thus reducing the natural rate of employment and benefit dependency. Also lower effective tax rates produced by a closer actuarial link between contributions and benefits may reduce unemployment and thus stimulate employment through this channel. To illustrate, Daveri and Tabellini (2000) find that the rise of 10 percentage points in the rate of effective labour tax in continental Europe in the seventies and eighties can explain about 3 percentage points of the increase in European unemployment during this period. Nickell and Layard (1999) estimate an unemployment effect of about 2 percentage points of such a tax increase.

[13] In the same vein, less parental leave is likely to raise employment.

dynamic revenue effect allows for a quantification of (some of) the efficiency gains from the introduction of IAs. The dynamic revenue effect 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 dynamic revenue gain reflects the efficiency gain from the increase in labour supply. It is given by the sum of the terms involving the labour-supply elasticities ε, η and φ in formula (4.6). The further revenue gain from reduced moral hazard – represented by the term $\chi c_m \left(\frac{m}{e}\right)$ in (4.6) – also implies a welfare gain, since the higher revenue allows the government to make (some) citizens better off.

In the presence of non-fiscal external effects, the overall welfare effect is comprised of more terms than the dynamic revenue effects alone. To illustrate, in case fertility produces positive non-fiscal external effects, the tax wedge for child benefits $c_m \left(\frac{m}{e}\right)$ should be amended by a term representing the net external effect from a child ψ, so that the overall first-order welfare effect of fewer children as a result of a higher value for α_m becomes $c_m \left(\frac{m}{e}\right) - \psi$. If child benefits are Pigovian (in the sense that child benefits are set equal to the non-fiscal external benefits so that $c_m \left(\frac{m}{e}\right) = \psi$), the 'co-insurance moral hazard effect' drops out from the expression for the welfare effects. Another example of a non-fiscal external effect is the expansion of employment in case of involuntary unemployment due to, for example, the market power of unions (see e.g. Daveri and Tabellini (2000)). In that case, the after-tax wage overstates the cost of employment and 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.

4.2. Illustration: An IA reform proposal for Denmark

To evaluate the likelihood that a switch to individual accounts will generate a Pareto improvement, we now consider a specific IA reform proposal for Denmark. There are four reasons why the introduction of IAs is likely to improve the equity-efficiency trade-off in a country like Denmark. First, the Danish system of social insurance is of the Beveridgean type, with a very weak link between taxes paid and benefits received. The bulk of social-insurance benefits is financed out of general tax revenues, and most benefits are paid out in flat rates unrelated to previous wages. Hence, the existing labour-income taxes financing intrapersonal redistribution over the life cycle incorporate a large distortionary element. Second, by international standards the effective tax and benefit rates tend to be high in Denmark, so the efficiency gains from cuts in these effective rates are potentially large. Third, as documented in Sørensen et al. (2006), the current Danish welfare state arrangements involve a large element of intrapersonal redistribution. Finally, compared to other countries, heterogeneity in gross (i.e. before-tax) lifetime incomes is only limited in Denmark.

Against this background, the Danish Economic Council (2005, ch. VI) has proposed that part of the existing social-insurance benefits be financed via mandatory individual savings accounts. According to this reform proposal (henceforth 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 an imputed labour income that 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.[14] 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

[14] When 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 below assume that all positive IA balances are paid out as lump sums at the time of retirement.

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. Grants to students in higher education
3. Short-term unemployment benefits
   (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 3 compares the distributional characteristics of these programs to those of some other important transfer programs.

(Table 3 about here)

The last column in Table 3 (based on the same data as those underlying Table 1) shows that early retirement benefits, education benefits, parental leave benefits and child benefits imply a low degree of lifetime-income redistribution from rich to poor. The Danish government grants to students in higher education are offered to *all* students without any means testing. This explains why this program is not very redistributive. Programs such as social assistance, disability benefits and housing benefits were excluded from the IA system, since they involve a high degree of redistribution of lifetime incomes, as indicated in the second column of Table 1. Furthermore, the degree of lifetime-income redistribution implied by benefits paid to workers suffering long unemployment spells (exceeding three months) is about 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, benefits paid during long sickness spells tend to be 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.

As an integral part of the estimation of disposable lifetime incomes reported in section 2, we have estimated the total transfers received by our representative sample of the Danish population over their working lives. From these estimates we have calculated the magnitudes $\alpha_e c_e \left(\frac{1-e}{e}\right)$ and $\alpha_e^* c_e^* \left(\frac{1-e^*}{e^*}\right)$ (or $\alpha_m c_m \left(\frac{m}{e}\right)$ and $\alpha_m^* c_m^* \left(\frac{m^*}{e^*}\right)$, depending on the type of transfer program) appearing in formula (4.6) in Box 4 – that is, the average amount of benefits withdrawn from the IA relative to average pre-tax income for individuals with an IA surplus and for the population as a whole, respectively. As explained in section 4.1, the difference between these magnitudes determines the size of the static revenue loss from an IA system, before accounting for the budgetary impact of behavioural changes. The first two columns in Table 3 show our estimates of $\alpha_e c_e \left(\frac{1-e}{e}\right)$ and $\alpha_e^* c_e^* \left(\frac{1-e^*}{e^*}\right)$ (and $\alpha_m c_m \left(\frac{m}{e}\right)$ and $\alpha_m^* c_m^* \left(\frac{m^*}{e^*}\right)$) The differences in these magnitudes between people with an IA surplus and the population as a whole depend partly on the size of the transfer program and partly on the skewness of the distribution of the benefits from the program. The differences between the numbers in the first and second columns in Table 3 also depend on the specific transfers included in the IA system, because this will determine the separation between people with and without an IA surplus. Furthermore, recall from expression (4.5) in Box 4 that the magnitudes $\alpha_e^* c_e^* \left(\frac{1-e^*}{e^*}\right)$ and $\alpha_m^* c_m^* \left(\frac{m^*}{e^*}\right)$ in the second column of Table 3 indicate the cut in the payroll tax rate made possible by including the relevant transfer program in the IA system.

To estimate the co-insurance participation effects captured by the bottom line of formula (4.6), we also need estimates for α_e and α_m, i.e., the fraction of total transfers to people of working age accounted for by the various transfer programs included in the IA system. These estimates are reported in the third column of Table 3. They were derived from the figures in the first column of the table. For example, in the case of early retirement benefits, the value of α_e was calculated by dividing the number in the top row of the first column by the number in the seventh row of that column (i.e., $\alpha_e = 0.0209/0.0509 = 0.4126$).[15]

The DEC proposed that the mandatory contributions to the IAs should be a fixed percentage of the base for the proportional Danish payroll tax ('arbejdsmarkedsbidrag'). This tax is levied on gross wage income and on the imputed labour income of the self-employed (with no cap for any of these groups); for wage earners, the tax is collected at the employee level. The revenue collected by this tax supplements general government revenues and is not earmarked for financing social insurance benefits. According to the DEC proposal, the percentage IA contribution would be set such that total contributions would correspond to total expenditure on the transfers included in the IA system, and the payroll tax rate would be cut by a corresponding amount, in accordance with the assumption underlying our formula (4.5) in Box 4. On the basis of this formula plus the figures in the second column of Table 3, we estimate that the proposed IA system covering the six transfer programs mentioned above would imply an 8.17 percentage-point cut in the payroll tax rate (and the introduction of a corresponding IA contribution). This is very close to the 7.9 percentage-points cut in the paytoll tax estimated by the DEC.[16]

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

[15] Combined with our assumption in section 4.4 that the replacement rate c_e is roughly the same for the various out-of-work benefits, this procedure implies that the non-employment rate $1 - e$ is assumed to be the same for the different benefit programs. This short-cut was necessitated by the fact that the dataset described in section 2 does not allow separate program-specific estimates of c_e and e (or of c_m and m). Although our rough assumption that $1 - e$ is the same across transfer programs undoubtedly implies some inaccuracy in our estimates of the *program-specific* co-insurance participation effects of IAs, it does not yield a systematic bias in our estimate of the *total* co-insurance effect of the entire IA system, given that the replacement rate is roughly the same across programs.

[16] The slight difference between the two estimates may be explained by the fact that the DEC would allow the IA contribution to be deducted against the base for the personal income tax, as is the case for the current Danish 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. If the marginal tax rate at the time of contribution to the IA differs from the marginal tax rate on payments from the account, these tax rules imply a slight modification of our formulas (4.5) and (4.6). However, as shown in the appendix, these modifications will only be of second-order magnitude if there is no IA system from the outset. This explains why our estimate of the payroll tax cut is so close to the estimate produced by the DEC.

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.

4.3. Effects on income distribution

Table 4 shows the estimated distributional impact of the proposed IA system, assuming a zero growth-adjusted real interest rate. The estimate was produced using the data and the constructed synthetic life cycles described in section 2.1. Importantly, the table abstracts from any behavioural effects of the IA system. The numbers thus reflect only the static impact effect. Although the very purpose of the IA system is to influence labour-market behaviour, the distribution of positive IA balances in Table 4 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 behaviour in the initial equilibrium and are not rationed on the labour market.[17]

(Table 4 about here)

The second row in Table 4 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 systematically rises with lifetime income. Moreover, the ratio of the average positive IA balance to lifetime income also rises with the income level, as shown in the third row in Table 4. 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 behaviour), 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

[17] The IAs do not have first-order implications for the welfare impact of capital-market imperfections such as liquidity constraints. The reason is that the IAs allow individuals to enjoy the same insurance benefits as under the current system – even if their account balance is negative. In this respect, the accounts provide the same liquidity insurance as current public benefits.

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, it would only rise to 0.133 if the proposed IA system were introduced (Danish Economic Council, 2005, ch. VI). 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.[18] Moreover, as we shall argue in the next section, the proposed IA system would generate a Pareto improvement even under rather conservative assumptions regarding behavioural responses.

4.4. Effects on the public budget and on economic efficiency

We will now employ formula (4.6) to estimate the revenue and welfare effects of the DEC proposal. When applying the formula, it must be recalled that the effective tax rates on labour income include indirect taxes; therefore, prior to the adjustment for a possible link between taxes and benefits, the effective marginal and average tax rates on labour income are given as

$$\tau = \frac{\tau^d + t^c}{1 + t^c}, \qquad t = \frac{t^d + t^c}{1 + t^c}, \tag{2}$$

where τ^d and t^d denote the marginal and average direct tax rates (before taking into account the additional benefit rights that might be generated by more employment and hours worked), respectively, and t^c stands for the overall effective indirect tax rate on consumption. In Denmark, realistic values of these tax parameters for an average worker

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

would be[19]

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

Another issue in the application of formula (4.6) is the estimation of our parameters ϕ_e and ϕ_h quantifying the degree to which increases in employment and hours worked generate additional benefit rights (in all transfer programs – not just in those included in the IA system). Since Danish public retirement benefits are unrelated to previous earnings, the discussion below refers only to transfer programs for people of working age. Suppose the benefit rate obtainable in some transfer program depends on previous earnings. In that case, a unit increase in earnings today will increase the future benefit rate by the replacement rate c_e, defined as the *after-tax* benefit rate relative to the *pre-tax* wage rate. Suppose further that, on average, the wage earner expects to be eligible for the benefit during some fraction u of his remaining labour-market career. With a zero growth-adjusted discount rate, this person's effective tax rate on labour income should then be reduced by the amount uc_e to reflect the fact that additional earnings generate additional future benefits. However, it may be the case that for some people eligible for the benefit, the benefit rate is unrelated to previous earnings, for example, because benefits are capped. Hence, we estimate our parameter ϕ_h by the simple formula

$$\tau\phi_h = \tau - a_h u c_e, \tag{3}$$

where a_h is the fraction of workers who are in a position to raise their future benefits by increasing their current working hours. Note that, in principle, the parameters a_h, u and c_e are averages across all transfer programs for those individuals of working age who (expect to) end up with a surplus on their IA.

In the Danish systems of unemployment insurance and sickness benefits, the benefit rate does in principle vary in proportion to previous earnings, but there is a relatively low cap on the benefit rates implying that only about 10 percent of full-time workers actually experience a direct link between their benefits and their previous wage income. However,

[19] 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.

for some part-time workers a move from part-time to full-time employment may cause an increase in the rates of unemployment and sickness benefits that they may collect if they become unemployed or get sick. In recent years the share of Danish workers working less than 30 hours per week has been slightly less than 9 percent. Persons working longer than that will generally not obtain additional benefit rights by increasing their hours worked. Against this background, we set $a_h = 0.2$ to cover that fraction of the workforce where a link between hours worked and future benefits received may reasonably be expected. This estimate will tend to understate the existing tax distortions (and hence to stack the deck against our IA proposal) because several transfer programs offer purely flat rates of benefit with no link between benefits and earnings for any benefit recipient. Furthermore, as a rough average across Danish transfer programs, the replacement rate c_e^* for the average worker would be around 0.25 (see Sørensen et al., 2006, Technical Appendix), but since people with an IA surplus are likely to have lower replacement rates because they tend to earn higher wages than the average worker, we set $c_e = 0.2$. By definition, the parameter u equals the non-employment rate $1 - e$. From the restriction that $\sum \alpha_e = 1$ across all programs offering out-of-work benefits plus the assumption that $c_e = 0.2$, the estimate in the seventh row of the first column of Table 4 implies that $u \approx 0.2$.[20] For comparison, the fraction of Danes of working age receiving some kind of public transfer has tended to hover around 0.25 in recent years. Our estimate of the non-employment rate for people with an IA surplus thus has the plausible implication that these individuals tend to be less dependent on public transfers than the average worker. With these parameter values we obtain $\phi_h = 0.987$. Thus we see that, for realistic parameter values in a Danish context, the average link between hours worked and future benefits received is so weak that our parameter ϕ_h will be close to unity.

[20] Here we make the simplifying assumptions mentioned in footnote 13 that e and c_e are the same across benefit programs (the assumption of roughly identical replacement rates c_e in the various out-of-work benefit programs is not a bad approximation in the Danish context). From the seventh row of the first column of Table 4 and the fact that $\sum \alpha_e = 1$ and $u \equiv 1 - e$, we then have

$$\sum \alpha_e c_e \left(\frac{1-e}{e}\right) = c_e \left(\frac{1-e}{e}\right) \sum \alpha_e = c_e \left(\frac{u}{1-u}\right) = 0.0509$$
$$\implies u = 0.203.$$

By analogy to (3), we estimate ϕ_e by the formula

$$t\phi_e = t - a_e u c_e, \tag{4}$$

where a_e is the fraction of people in the workforce who can increase their future benefit rights by moving from non-employment into employment. In Denmark, a member of an unemployment insurance scheme is entitled to unemployment insurance benefits if he/she has been employed for at least one year during the latest three years. Uninsured workers are entitled to a means-tested social assistance benefit in case of unemployment, provided they are available for work. The benefit depends on the current income and wealth of the household, but not on the previous earnings of the recipient. This means test may imply that also the spouse faces a disincentive to participate in the labour market, thus tending to increase rather than decrease our parameter ϕ_e. To be entitled to sickness benefits, a person has to have been employed only during the previous eight weeks. Entitlement to early retirement benefits is obtained by 25 years of membership of an unemployment insurance scheme, but members above the age of 50 can effectively go unemployed for as long as ten years without losing entitlement to early retirement benefits from the age of 60. Most other transfer programs involve no link between employment and benefit entitlement. Estimating the 'correct' value of a_e on the basis of these benefit eligibility rules is quite difficult, but an upper bound on a_e would certainly be given by the proportion of Danes of working age (aged 18-65) who are non-employed. As already mentioned, this fraction is roughly 25 percent; we therefore rather conservatively insert $a_e = 0.2$ along with our earlier estimates $u = 0.20$ and $c_e = 0.2$ into formula (4) to obtain $\phi_e = 0.985$. Again, we thus conclude that in the Danish labour market the non-distortionary fraction of the labour income tax is very low indeed.

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 we conservatively set $\eta = 0.1$. This relatively low estimate partly reflects the effectiveness of Danish active labour-market policies in encouraging transfer recipients to find work. Furthermore, by selecting low values for the labour-supply elasticities, we account for the possibility that some agents may be myopic or liquidity constrained and therefore do not fully account for the intertemporal links between withdrawals from the accounts and future retirement benefits.

Finally, formula (4.6) requires a distinction between those benefits in the IA system that are directly related to the recipient's employment status and those that are not. In the latter category we include education benefits and child benefits, whereas unemployment benefits, sickness benefits, early retirement benefits and parental leave benefits are clearly paid out only during periods of non-employment.

For the latter four transfer programs, we can now estimate the effects of introducing IAs on the basis of the assumptions and information stated above, since the parameters χ and φ in formula (4.6) are not relevant for these programs (in terms of the formula, $\alpha_m = 0$ for these programs). Inserting our assumptions on parameter values plus the relevant estimates from the three first columns of Table 3 (which, as indicated above, are derived from our estimates of lifetime incomes in section 2) into formula (4.6), we obtain the estimated budgetary effects stated in the five top rows of Table 5.

(Table 5 about here)

Given our assumed parameter values, including these four programs in an IA system would improve the public budget by more than three percent of the initial labour income tax base for those people who end up with an IA surplus. According to the data underlying Table 5, ewh amounts to almost 60 percent of the total labour income tax base, so the estimated gain in net revenue is almost two percent of the total tax base. Table 5 shows that the bulk of the dynamic net revenue gain comes from the participation response to the cut in effective benefit rates implied by the IA system (see column 4 in the table). This is not surprising, considering that the IA system effectively cuts the replacement rate by 100 percent in those transfer programs that are included in the system.

To illustrate the workings of the account system and to explain how the numbers in Table 5 were calculated, it may be instructive to consider the first row in Table 5 which shows the various effects of including short-term unemployment benefits in the IA system. The inclusion of this program in the account system allows a cut in the marginal labour-income tax rate of about 2.12 percentage points, as indicated in the fourth row and second column of Table 3. Multiplying this tax rate cut by the factor $\varepsilon\left(\frac{\phi_h \tau}{1-\phi_h \tau}\right)\bar{\phi}_h$

appearing in formula (4.6), we can estimate the rise in tax revenue generated by the increase in working hours resulting from including short-term unemployment benefits in the IA system. Clearly, this effect depends on the hours-of-work elasticity ε and the initial level of the marginal effective labour-income tax rate, $\phi\tau$. Using the parameter values mentioned above, we estimate an increase in tax revenue amounting to 0.18 percent of the labour-income tax base (for individuals with an IA surplus), as shown in the first row and second column of Table 5.[21]

The *average* tax rate on labour income will also drop by 2.12 percentage points when short-term unemployment benefits are included in the IA system because the Danish payroll tax is purely proportional. This will stimulate labour supply at the *extensive* margin, as the lower average tax burden on labour income induces the unemployed to increase their search efforts in order to move more rapidly into employment. The resulting effect on the public budget is captured by the term $\left(\frac{\eta(\phi_e t+c_e)}{1-\phi_e t-c_e}\right)\bar{\phi}_e$ in formula (4.6). This term includes the 'participation elasticity' η (in this case picking up the effect of more intensive job search) and the initial effective 'participation tax rate' $\phi t+c_e$, which reflects the increased tax burden and the loss of benefit income experienced by an individual who moves from unemployment into employment. Given our assumptions on parameter values, we estimate that the cut in the average effective labour income tax rate increases the employment rate by about 0.78 percent. This will in turn improve the budget by 0.57 percent of the labour income tax base, as reported in the first row and third column of Table 5.

Finally, since collection of short-term unemployment benefits reduces a worker's IA balance by a similar amount in present-value terms, the inclusion of these benefits in the IA system effectively reduces the replacement rate to zero. This provides a further strong incentive for forward-looking individuals to raise their labour supply at the extensive margin through increased search effort. The resulting effect on the budget is reflected in the term $(\phi_e t + c_e)\left(\frac{c_e\eta}{1-\phi_e t-c_e}\right)\alpha_e$ in formula (4.6), where the estimate for α_e is given in the fourth row and third column of Table 3. On this basis, we obtain a budgetary improvement of 1 percent of the labour income tax base, as stated in the first row and fourth column of Table 5. The total dynamic effect on the budget is the sum of the

[21] In applying formula (4.6), we use the facts that $\overline{\phi}_h \equiv d(\phi_h\tau)/d\tau = \phi_h$ and $\overline{\phi}_e \equiv d(\phi_e\tau)/d\tau = \phi_e$.

three effects just mentioned, adding up to 1.75 percent of the labour-income tax base of individuals with a positive IA balance (who account for 60 percent of the total labour-income tax base).

The numbers in the other rows in Table 5 were estimated in a similar manner. An important caveat should be mentioned. If people are liquidity-constrained and myopic, the mandatory IA contribution will work to some extent like an ordinary tax; for these individuals the cut in the effective tax burden will be lower than indicated in Table 3. Moreover, these agents will not (fully) account for the impact on their IA balance of taking up some benefit. The improvement in labour-market incentives may thus be smaller than our calculations assume. We accounted for this bias in a pragmatic way by adopting rather conservative assumptions on the magnitude of labour-supply elasticities.

In the absence of non-fiscal external effects, the dynamic revenue gain is a proxy for the welfare gain from the introduction of IAs, as we explained in section 4.1. From column 4 in Table 5, this efficiency gain can estimated to be a respectable 3.5 percent of the *total* labour-income tax base ($0.6 \cdot 5.81 \approx 0.035$), despite our rather cautious assumptions on labour-supply elasticities. Note that since the proposed IA system ensures that nobody can be financially worse off than under the existing tax-transfer system, the estimated efficiency gain represents a genuine ex-post Pareto improvement, provided the four transfer programs do not serve to correct for non-fiscal externalities. One might argue that the parental leave scheme has a positive external effect in so far as parental care in the early stage of childhood improves the social skills of the child. If there are negative non-fiscal externalities associated with non-employment (e.g., increased crime rates and loss of social skills), however, the increase in employment obtained through the IA system would have positive external effects that are not included in Table 5. Thus it is not at all obvious that an allowance for non-fiscal externalities would reduce the estimated total welfare gain. At any rate, the estimate in the fifth row of Table 5 shows how large the possible loss in non-fiscal externalities would have to be to eliminate the total net welfare gain from the IA system.

Estimating the revenue- and welfare effects of incorporating education benefits and child benefits in the IA system is a more complex matter, requiring an estimate of the parameters χ and φ in formula (4.6). In the case of grants to students in higher education, the elasticity χ measures how much the fall in the effective benefit rate reduces the number of such students. Unfortunately, little is known about the size of this elasticity. Maintaining our focus on education grants, the elasticity φ reflects how a change in the number of students in higher education affects total employment. Here there are offsetting effects. When the rise in the private cost of higher education induces some (potential) students to join the labour force at an earlier age, this will *ceteris paribus* stimulate employment. But with a lower average level of education, there is an offsetting negative effect on the size of the labour force measured in efficiency units. Moreover, since higher-educated people tend to retire later, a fall in their number will also tend to reduce the average retirement age. A priori, the sign of φ (not to mention its size) is thus unclear.

In the case of child benefits, χ measures how their inclusion in the IA system affects fertility. Presumably, the cut in the effective benefit rate would reduce fertility, but little is known about the likely magnitude of this effect. The effect of lower fertility on the rate of employment is captured by our parameter φ. A lower fertility rate would tend to raise the number of women in the labour force, but it might also reduce the pressure on males to live up to their traditional role as breadwinners in families with dependent children, causing a drop in male labour supply. Thus, even though the net effect on labour supply would probably be positive, the quantitative effect seems highly uncertain.

In itself, the likely fall in the number of students and in the number of children will reduce government expenditure. This effect is captured by the last term $\chi c_m \left(\frac{m}{e}\right)$ in the bottom line of formula (4.6). The fiscal externality arising from the employment effect of the changing number of students and children is reflected in the second-last term $\chi\varphi\left(\phi_e t + c_e\right)$ in (4.6). At least in the case of child benefits, we argued that this impact on the public budget is also positive. The positive fiscal externality of the increased labour supply stemming from lower effective tax rates is represented by the first term $\left[\varepsilon\left(\frac{\phi_h \tau}{1-\phi_h \tau}\right)\bar{\phi}_h + \left(\frac{\eta(\phi_e t + c_e)}{1-\phi_e t - c_e}\right)\bar{\phi}_e\right] c_m^* \left(\frac{m^*}{e^*}\right)$ in the bottom line of (4.6). In the absence of non-fiscal external effects from higher education and population growth, the sum of these three terms would represent the welfare gain from including education grants and child benefits in the IA system, provided that net government revenue does not go down.

In the benchmark case where the initial benefit rates have been chosen to ensure that the fiscal externalities from changes in the number of children or students are roughly offset by the non-fiscal and direct fiscal externalities, one may as a first approximation ignore the contributions of the terms $\chi\varphi\left(\phi_e t + c_e\right)$ and $\chi c_m\left(\frac{m}{e}\right)$ to the welfare effect of the IAs. This is obviously convenient, given that so little is known about the parameters χ and φ. In this case, the welfare effect can be evaluated from the first three columns in the two bottom rows of Table 5. According to the estimates presented there, the inclusion of education benefits and child benefits would be roughly revenue neutral, ensuring that the net taxes on people with an IA deficit would not have to be raised. At the same time, persons with an IA surplus would gain, implying a Pareto improvement. In this scenario the size of the aggregate efficiency gain may be approximated by the sum of the figures in columns 2 and 3.

It might be argued that when it sets the pre-reform benefit rates, a rational government would not only account for the fiscal and non-fiscal externalities from changes in the number of children and students; it would also account for the indirect fiscal externalities in columns 2 and 3 which stem from the need to finance benefits through distortionary taxes. Hence, it would seem that the introduction of IAs could not generate a Pareto improvement if the government has already optimized the tax-transfer system. However, this argument overlooks the fact that the introduction of IAs provides the government with a new policy instrument that allows a selective cut in taxes and benefits for those who manage to accumulate an IA surplus, thereby reducing the distortions arising from intrapersonal redistribution. With such an additional policy instrument at its disposal, the government should be able to make some people better off without making anybody worse off, even if it has made optimal use of all of its previous policy instruments.

The numerical exercises in this section seem to support our earlier conjecture that the introduction of IAs for those transfer programs that involve relatively little interpersonal redistribution could potentially be Pareto-improving, provided the IA system includes a 'lifetime income guarantee' (a bail-out clause) for those who end up with a negative account balance.

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 labour-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 labour-market institutions that combat long-term employment and facilitate rapid turnover and by social policies that provide a lifetime 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 labour-market incentives, individual accounts with a lifetime income guarantee seem to be an attractive alternative to simple cuts in taxes 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, individual accounts 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 labour-market policies more on the lifetime poor, thereby also protecting the social norms of this group.

We have explored whether the introduction of individual accounts for certain public transfers can yield a Pareto improvement. The optimal size of individual accounts is a more complicated issue that we leave for future research. Individual accounts also have implications that have not yet 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 (see, e.g., Casamatta, Cremer, and Pestieau (2000)). 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 labour market. The lifetime income guarantee implies also that, while marginal rates are cut for others, marginal tax rates tend to remain large only for the lifetime poor.[22] The employment gap between low-skilled and high-skilled workers may thus increase

[22] A 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.

unless 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 lifetime-income guarantee. If the government uses 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 labour-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 combatting moral hazard in social insurance. Through all of these channels, individual accounts support social policy by reducing the costs that are associated with an effective mix of redistribution, social insurance and consumption smoothing.

REFERENCES

Arrow, K.J. (1971). *Essays in the Theory of Risk Bearing.* Chicago: Markham Publishing Co.

Belan, P., and P. Pestieau. (1999). Privatizing social security: a critical assessment. Geneva Papers on Risk and Insurance – Issues and Practice 24, 114-130.

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. M.I. Hansen and P.B. Sørensen (2006). Estimating lifetime incomes. Technical appendix downloadable at www.econ.ku.dk/pbs/.

Brown, A.J.G., J.M. Orszag and D.J. Snower, 2006, Unemployment accounts and employment incentives. Kiel Working Paper No. 1274.

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 labour, 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.

Cassamatta, G., H. Cremer, and P. Pestieau (2000). Political sustainability and the design of social insurance. Journal of Public Economics 75, 341-364.

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 Incitamenter 2002* (Distribution and incentives 2002). Copenhagen.

Danish Ministry of Finance (2004) *Fordeling og Incitamenter 2004* (Distribution and incentives 2004). Copenhagen.

Daveri, F., and G. Tabellini (2000). Unemployment and Taxes: Do Taxes Affect the Rate of Unemployment? *Economic Policy,* pp. 49-104.

Falkingham, J., J. Hills, and C. Lessof (1993). William Beveridge versus Robin Hood: Social security and redistribution over the life cycle. WSP/88. STICERD, London.

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.

Gollier, C. and H. Schlesinger (1995). Second-Best Insurance Contract Design in an Incomplete Market. *Scandinavian Journal of Economics* No. 97, pp. 123-135.

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 - Velfærdsydelser finansieret af individuelle opsparingskonti (Citizen Accounts - Social insurance financed by individual savings accounts). Master thesis, Department of Economics, University of Copenhagen.

Harding, A. (1993). *Lifetime Income Distribution and Redistribution.* North-Holland, Amsterdam.

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.

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. Forthcoming 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.

Nickell, S., and R. Layard (1999). Labor Market Institutions and Economic Performance, in O. Ashenfelter and D. Card (ed.), *Handbook of Labor Economics*, Vol. 3 (North Holland, Amsterdam).

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.

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.

Sinn, H.W. (2000). Why a funded pension system is useful and why it is not useful. International Tax and Public Finance 7, 389-410.

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.

Sørensen, P.B., M. Hansen and A.L. Bovenberg (2006). Individual savings accounts and the life-cycle approach to social insurance. Paper presented at the Netspar conference on *Reinventing the Welfare State* in The Hague, April 27-28 2006.