How Much Intergenerational Risk Sharing Does the U.S. Social Security System Really Provide?

Libor Dušek*

June 2011

Abstract

The paper documents the intergenerational risk sharing that the U.S. Social Security system has actually provided throughout the 20th century. It investigates the relationship between lifetime net transfers that the cohorts born between 1900 and 1985 expected to receive from Social Security and their lifetime wages and returns on savings. Differences between cohorts in lifetime net transfers are negatively related to the differences in returns to savings but not to differences in lifetime wages. Aggregate shocks to wages translate into reductions in the net transfer to the young and an increase in the net transfer to the old; the exact opposite pattern holds for aggregate shocks to returns on savings. A large fraction of the differences between cohorts in lifetime net transfers cannot be explained by risk sharing considerations.

Keywords: social security, intergenerational risk sharing, policy risk, pension reform

JEL Codes: H55, D72

*CERGE-EI, Prague (a joint workplace of the Center for Economic Research and Graduate Education, Charles University, and the Economics Institute of the Academy of Sciences of the Czech Republic), Politickych veznu 7, Praha 1, Czech Republic. libor.dusek@cerge-ei.cz, tel. +420-224-005-146, fax. +420-224-005-333. The research was supported by a Global Development Network grant RRC X-67/2010. Research assistance from Martin Kuncl and Jiří Šatava, and comments from Randall Filer, Vincenzo Galasso, Filip Matějka, Sita Slavov, and participants of the WDA/scala workshop on Pension Challenges and the Financial Crisis in St.Gallen and Prague Economic Meeting are kindly appreciated. I am particularly grateful to Sita Slavov for sharing her data set.

1 Introduction

The ability to share risk across generations is a desirable feature of pay-as-you-go pension systems. A generation that suffers a negative economic shock (such as a stock market crash just before retirement or a major recession during working years) can in principle be compensated by higher benefits or lower contributions. The compensation can be provided by the relatively better-off generations that are alive concurrently with the affected generation or will be alive in the future.

This paper assesses whether the Social Security system as it has been implemented in the United States has provided intergenerational risk sharing in this manner. Specifically, it investigates whether cohorts that earned relatively lower incomes in the market received relatively larger net transfers from the Social Security over their lifetimes. The main contribution is in investigating the Social Security "as actually implemented" since the 1930's; I document the relationship between incomes and net transfers that the cohorts actually received from the Social Security based on the continuously evolving legislation that prescribed the contribution rates and benefits.

This approach differs from the approach taken by the existing literature on the intergenerational risk sharing properties of the pay-as-you-go pension systems. One line of the literature (Gordon and Varian 1988, Rengel and Zeckhauser 2001, Ball and Mankiw 2007 or Bohn 2009) builds overlapping generations models in which incomes are subject to cohort-specific shocks and markets fail to provide ex-ante efficient allocation of consumption across generations. A properly structured scheme of intergenerational transfers can allocate consumption across generations in a Pareto-improving manner. While the papers vary in their assumptions, they broadly agree on qualitative prescriptions. For example, if a particular generation suffers a negative income shock, all current and future generations should provide a transfer to the affected generation (if such transfers are technologically feasible); the compensating transfers should be spread equally across the other generations; if the young suffer a negative shock, the old should also provide a transfer (i.e., the transfers need not necessarily flow from the young to the old).[1]

A closely related literature (Wagener 2004, Auerbach and Lee 2009, Gottardi and Kubler 2011)

[1] The models differ, though, in quantitative predictions, such as whether a shock to a particular generation should result in an equal change in income for all generations (complete risk sharing) or whether the affected generation bears a disproportionate share of the shock. Complete risk sharing comes from the Gordon and Varian (1988) basic model. Bohn (2009) implies that the elasticities of consumption with respect to any shock should be the same for the young and the old, that is equal proportionate change for all generations. In D'Amato and Galasso (2008) the optimal transfer between generations is a linear function of the shock (to the return on savings) and in general less than proportional.

evaluates the risk sharing properties of stylized pension systems. These papers assume a stochastic process for economic shocks, simulate the distribution of possible histories of economic outcomes, compute how a given pension system translates the outcomes into contributions, benefits, and the resulting consumptions under each possible history, and calculate the ex-ante expected utilities. The pension systems are explicitly treated as "stylized" - the rules stipulating the contribution rate, the benefit formula, or the retirement age are assumed to be in place forever. Such analysis is useful for the optimal design of the pay-as-you-go pension systems.

In reality, the rules change. Politicians adopt pension reforms now and then; sometimes minor, sometimes radical. As a result, each cohort lives through a series of reforms, each of them changing the stream of future contributions and benefits that the cohort is being promised. During the 66 years of Social Security covered in this paper, there were 21 legislative changes that had a substantive effect on the value of contributions and benefits.[2] The resulting intergenerational transfers, "as implemented" by ever-changing rules, may be very different from the transfers that would emerge if any of the rules was in place forever. Several political economy models actually predict that the politicians implement pension systems that do not share risk across generations optimally: In D'Amato and Galasso (2010), the political equilibrium provides transfers to the old that are too stable. In Rengel and Zeckhauser (2001), transfers from the young to the old are much more likely to constitute a political equilibrium than reverse transfers even though reverse transfers are often times necessary for optimal risk sharing.

Understanding the risk sharing properties of "as implemented" pay-as-you-go pension systems is potentially useful for making long-term projections of the pay-as-you-go system's finances or for evaluating the optimal balance between the funded and the pay-as-you-go pillars. The literature has so far documented that certain cohorts experienced substantial negative shocks due to selected pension reforms. McHale (2001) calculates the changes in social security wealth[3] induced by pension reforms implemented in the G7 countries during the 1990's for workers with average earnings at age forty-five and at the standard retirement age. He finds that some of the reforms reduced social security wealth by as much as 29% (the Italian 1992 reform) or 26% (the German 1992 reform). He also finds that those at the retirement age experienced only minor, and in most cases no, cuts in benefits. Borgmann and Heidler (2003) compute the changes in relative generosity of the German pension system and show that they are largely influenced by changes in the population projections. Dusek and Kopecsni (2008) compute changes in the social security wealth for different cohorts, genders, and education levels due to all major pension reforms in

[2]Shoven and Slavov (2006).

[3]Social security wealth is defined as the expected present value of social security benefits minus the expected present value of social security taxes from a given point in time forward.

Hungary, Czech Republic, and Slovakia. The authors document that some of the reforms rather created additional intergenerational risk by inducing rather arbitrary changes in social security across cohorts.

Leimer (1994, 2007) investigated intergenerational transfers throughout the Social Security's history by calculating the lifetime net transfers from Social Security for all cohorts affected by the program. The reported transfers vary substantially by cohort. In a paper that is closest to ours, Shoven and Slavov (2006) compute how the lifetime net transfers varied over time at the cohort level. They calculate the internal rate of return on contributions in the U.S. Social Security system for cohorts born in 1900, 1905, 1910, etc till 1985 in each year based on the legislation valid in that year. The IRRs also varied substantially within a cohort over time; measures taken to restore the financial solvency of the system were the major direct cause of that variation.

The intergenerational transfers "as implemented" by the Social Security since 1939 till today are also the subject of this paper. The main contribution over Leimer (2007) and Shoven and Slavov (2006) is in investigating whether, and to what extent, the relative differences between cohorts in the lifetime net transfers can be explained by relative differences in the cohort's economic outcomes, namely lifetime wages and returns on savings. I propose and estimate simple policy functions to characterize the relationships between lifetime net transfers, wages, and returns on savings across cohorts and discuss whether the observed relationships are qualitatively consistent with optimal intergenerational risk sharing.

I use data on the histories of wages, returns on savings, contributions and benefits that a representative agent was expected to realize in each year from 1939 till 2005, based on the social security legislation valid in that year and on adaptive expectations about the future.[4] From those I construct the real present value of lifetime wages, excess savings (savings accumulated at the end on the agent's working history minus the savings that the agent would have accumulated if the returns on savings were constant) and net transfers (Social Security benefits received minus contributions paid). The data set covers cohorts born in 1900, 1905, 1910, etc. till 1985. I investigate the relationship between lifetime net transfers and economic outcomes at two levels. The first one is a cross-section of "terminal" outcomes for each cohort, i.e., the lifetime wages, excess savings, and net transfers that the cohort realized by the end of its life, or for younger cohorts, that they realized by the last year of the data set and expected to realize for the remainder of their life. The second one is a panel of lifetime outcomes that the cohorts

[4] The histories of expected wages, contributions, and benefits are taken from Shoven and Slavov (2006) data set. I constructed the histories of returns on savings.

expected every five years, i.e., in 1940, 1945, etc. till 2005. It allows analyzing how the expected net transfers adjusted with shocks to wages and returns on savings.

The cross-sectional comparison of terminal outcomes shows that differences in lifetime excess savings translate almost one-for-one into differences in lifetime net transfers for the pre-1960 cohorts, and more than one-for-one for the post-1960 cohorts. Social Security system thus appears to have provided complete (possibly more than complete) risk sharing of the stock market risk. This finding is particularly striking given the fact that the Social Security rules do not explicitly link contributions and benefits to the stock market returns.

The panel analysis of expected lifetime outcomes again finds some evidence of risk sharing with respect to excess savings although much smaller in magnitude than in the cross-section. Over the Social Security history, an increase in excess savings of one cohort by 1 percent of lifetime income is associated with a reduction in the net transfer by at most 0.15 percent of lifetime income, holding excess savings of other cohorts constant. The latter result is driven by comparing cohorts within a year; when comparing pairs of cohorts over time there is no evidence of risk sharing even with respect to excess savings. Neither the panel nor the cross-sectional analysis finds evidence of risk sharing with respect to wages.

I also estimate how aggregate shocks[5] are reflected in the changes in the lifetime net transfers of different cohorts. A 10-percent increase in aggregate lifetime wages is associated with a reduction in the net transfer by 0.19 percent of lifetime income for the youngest cohorts (aged below 20) and an increase in the net transfer by 0.34 percent of lifetime income for the already retired cohorts. The exact opposite pattern holds for aggregate savings - a 10-percent increase is associated with an increase in the net transfer by 0.27 percent of lifetime income for the youngest cohorts and a reduction in the net transfer by 0.35 percent of lifetime income for the already retired cohorts. These adjustments, while small in magnitude, are broadly consistent with optimal intergenerational risk sharing.

The observed relationships between net transfers and wages/excess savings are different for cohorts for which we observe completed (or near-completed) lifetime outcomes and for cohorts for which the lifetime outcomes are based mostly on expectations about the future. The Social Security system "as implemented" with numerous reforms accumulating during the cohorts' lifetime provides different pattern of risk sharing than if the system's default rules were held unchanged forever.

[5] An aggregate shock to lifetime wages is defined as the average change in expected lifetime wages over the 5-year period across all cohorts below the retirement age.

2 Data

Our key variables are expected lifetime wages, excess savings, and net transfers that a representative worker in each cohort expected to receive in each year. Precisely, they are defined as the present value of the realized past and expected future flows of real wages, excess savings, and net transfers. They are constructed from a data set created by Shoven and Slavov (2006)[6] who compute the internal rate of return on contributions in the U.S. Social Security system for each cohort born in 5-year intervals after 1900 (1900, 1905, 1910, till 1985) in each year between 1939 and 2005 under the social security legislation valid in that year. To do so, they first estimate the age-wage profile for an average worker and scale it by the average wage in each year to obtain the age-wage profile for each year. The worker's wage then evolves as he moves along the profile and the profile also shifts up with the average wage growth. They assume that the worker works from age 20 till the official retirement age and has a deterministic length of life of 80 years. For each year, they compute the contributions paid and benefits received historically based on the past social security rules, and construct the projections of future wages and inflation rates, based on adaptive expectations (the future growth rate of a variable will be equal to the average growth rate from the previous five years).[7] Finally, they construct the projections of future contributions and benefits from the wage and inflation projections by applying the social security legislation that was on the books in that year.

I construct the lifetime wages, excess savings, and net transfers by first converting the realized and projected wages etc. into real (2005) dollars. Next, I compute their present values based on a 4% discount factor. All variables are discounted to the year when the cohort is born, which implies that comparisons across cohorts are not affected by differential discounting.[8] Finally I construct the variable "lifetime net transfer" as the difference between the (present value of) real lifetime benefits and real lifetime contributions.

I also generate the evolution of expected lifetime savings for each cohort. Although the stock market shocks are common to all cohorts in a given year, the timing of the shocks during the lifecycle of the cohort generates variation in the amount of savings that each cohort accumulates upon retirement.[9] Based on Poterba et al. (1998) I assume that each cohort saves 9% of its

[6] I am extremely grateful to Sita Slavov for the willingness to share the data.

[7] All cohorts therefore experience the same shock to the level of their wages and expected wage growth rate. Still, the common shock has a differential impact on the lifetime wages of each cohort, since a younger cohort will reap an unexpected increase in wages for a longer period than an older cohort.

[8] For example, wages received in year 1940 by a 1900 cohort are discounted by the same discount factor as wages received in 1965 by a 1925 cohort.

[9] First, one cohort may experience higher average return during its working years than another cohort. Second,

annual wages towards retirement every working year.[10] The savings are invested in a portfolio consisting of 60% stocks that yield a real return equal to the inflation-adjusted return on the S&P500 index, and 40% in bonds that yield a safe 1% real return.[11] I construct the histories of accumulated and expected savings for each cohort and year by assigning the portfolio returns to the cohorts' accumulated savings. The stock returns expected in the future are assumed to be equal to the average S&P return during the past 20 years. I then compute the present value of (realized or expected) lifetime savings in each year, which is equal the to savings accumulated at the year of retirement (if the cohort has already retired by the year of expectation) or the savings that the cohort expects it will have accumulated at the year of retirement (if the cohort is still working). From that I subtract the savings that the cohort would have achieved if the return on savings were constant in all years and equal to the average return from 1901 till 2005. The difference gives the variable of interest referred to as ”lifetime excess savings”. The excess savings are used instead of the simple level of savings because savings are automatically higher for cohorts with higher wages, and thus higher savings for a particular cohort capture both a good stock market history or a good wage history. By removing the part of the savings that a cohort accumulated through higher wages, I isolate the former from the latter. In some calculations I also use lifetime income, which is the sum of the lifetime wages and the return on savings.

3 Relationship between net transfers and economic outcomes

This section documents the degree of intergenerational risk sharing through simple policy functions that relate differences between cohorts in lifetime net transfers to differences between cohorts in lifetime wages and excess savings. The policy functions are based on a simple idea that cohorts with relatively lower incomes should be compensated by relatively higher net transfers. They are estimated at two levels of observation: The first one evaluates only the ”terminal” outcomes, i.e. the present values of lifetime wages, savings, and net transfers observed in the final year of the cohort's life if the cohort has died by 2005, or, if the cohort is alive as of 2005, the outcomes expected observed in 2005. It is essentially a cross-sectional comparison between cohorts. The second exploits the panel structure of the data and evaluates the expected lifetime

a large shock (such as a stock market crash) has a larger affect if it hits a cohort shortly before retirement, since it changes the value of most their accumulated savings, while it may have little effect on the lifetime savings if it hits the cohort in its early working years.

[10]Poterba et al. (1998) reports that the average 401(k) contribution represents 9% of the contributing household's income.

[11]This division of assets corresponds to calculations in Feldstein and Ranguelova (2001).

transfers, wages, and excess savings as they evolved from 1939 to 2005.

3.1 Terminal outcomes

Figures 1 and 2 show scatter plots of lifetime net transfers against lifetime wages and lifetime excess savings, respectively. Labels on each data point denote the cohorts. The cohorts that were born later generally have higher wages, but it is not generally the case that they also received lower net transfers. There are groups of cohorts (1900-1920, 1935-1960) for which higher lifetime wages were associated with higher lifetime transfers. The cohorts born until 1945 received positive net transfers while all successive cohorts received negative transfers.[12] Only within the group of the youngest cohorts (1965-85) there is a negative relationship between transfers and wages, essentially dollar-for-dollar.

The relationship between lifetime net transfers and excess savings is radically different. It is uniformly negative[13], despite the fact that there is substantial variation in the excess savings by the birth year of the cohort (e.g., cohorts 1900, 1905, and 1935 experienced better-than-average stock market histories while the adjacent 1910 and 1930 cohorts experienced some of the worst stock market histories). There are again two distinct ranges of data: when comparing within the group of 1900-1955 cohorts, the net transfer decreases nearly dollar-for-dollar with an increase in excess savings, suggesting that the Social Security provided essentially complete intergenerational risk sharing against shocks to the returns on savings. This is particularly surprising given the fact that the Social Security rules do not contain explicit link between the stock market returns and the level of benefits. But "as implemented", the Social Security system appears to function as if it was perfectly compensating for between-cohort differences in returns on savings. Within the 1960-85 cohorts, the net transfer declines much faster with an increase in excess savings.

The clear structural break in both figures for the 1960-1985 cohorts has economic significance. The cohorts born before 1960 have realized their histories of wages, savings, contributions and benefits either fully (have already died) or almost fully (are already retired or near retirement). The computed lifetime net transfers reflect (fully or to large extent) the Social Security system "as implemented", with a long history of rules and reforms driving the resulting net transfers. The 1960-1985 cohorts have a fairly short working history and the computed lifetime net transfers group reflect (almost fully) the rules of Social Security system that were in force in 2005 and are

[12]This pattern of net transfers has been well documented (Leimer 2007).

[13]The 1935 cohort is an outlier as it experienced a relatively better stock market history yet received the highest net transfer of all cohorts.

assumed to be in place in the future. Those (younger) cohorts are expected to experience better wage as well as stock market histories and the current Social Security rules do generate lower net transfers to them. But the relationship between net transfers and wages or excess savings is very different for cohorts for which observations reflect the Social Security "as implemented". Risk sharing properties of the "stylized" and "as implemented" pay-as-you-go system appear very different.

A simple policy function (equation 1) formalizes the relationship between differences in the cohorts' lifetime wages and excess savings and differences in their lifetime net transfers. T, W, and S denote per capita lifetime net transfers, wages, and excess savings for each cohort. The subscripts i, k indicate a "comparison" cohort and a "benchmark" cohort. Differences between cohorts are normalized by the lifetime income of the "benchmark" cohort Y_k.

$$\frac{T_i - T_k}{Y_k} = \beta^W \frac{W_i - W_k}{Y_k} + \beta^S \frac{S_i - S_k}{Y_k} + \lambda_k + \epsilon_{ik} \tag{1}$$

If Social Security shares risk between cohorts, the coefficients β^W and β^S should be negative and between minus one and zero. Of course, factors other than the risk sharing considerations determine the differences in net transfers between cohorts. Consider three cohorts, i, j, and k which earned different incomes due to unexpected shocks such that the cohort i is the poorest and cohort k is the richest. The political process that determines the net transfers may take into consideration two factors: the political influence of each cohort and the desire to optimally share risk across generations. If only risk sharing considerations determine the transfers, then cohort i should receive the highest and cohort k the lowest transfer. If both political influence and risk sharing considerations determine the transfers and cohort k happens to be the most influential (such that it receives the highest net transfer despite being the richest), the risk sharing considerations vis-a-vis the remaining cohorts still imply that the gap in net transfers between cohorts i and k should be smaller than between j and k. The "benchmark" cohort dummy λ_k thus captures the fact that the differences in net transfers between any cohort i and cohort k would be systematically different (smaller).[14] The policy function 1 is estimated in a regression where the unit of observation is a pairwise combination of a "comparison" cohort i and all subsequent "benchmark" cohorts k, as if each ik observation was an observation of a subject k in different points in time in a conventional panel set-up.[15] The estimates are then identified out of deviations from the mean difference in net transfers, wages, and savings for a

[14]The coefficients on the cohort dummies can be interpreted as measures of such relative differences in influence.

[15]Only pairwise combinations where $i > k$ are included in the regression, since combinations with $k > i$ would only replicate the same observations but with the opposite signs.

"benchmark" cohort.[16]

Table 1 presents the estimates. According to the specification in column (1), an increase in the difference in lifetime wages between two cohorts by 1 percent of lifetime income actually reduces the difference in lifetime net transfers by 0.14 percent of lifetime income, while an increase in the difference in lifetime excess savings reduces the difference in lifetime net transfers by 2.03 percent of lifetime income. Column (2) presents an alternative specification where the "benchmark" cohort fixed effects are replaced by the "comparison" cohort fixed effects. The coefficient on the lifetime wages is now negative and the coefficient on lifetime excess savings remains negative but is smaller in magnitude (-0.600). The last column of Table 1 excludes the fixed effects in order to show how much of the variation in net transfers is explained by the cohort fixed effects. The R^2 is 0.45 when excluding the fixed effects which is by 0.15 or 0.5 lower compared to the specifications with fixed effects.[17] Large part of the variation in net transfers is thus unrelated to differences in lifetime incomes between cohorts. This finding can be interpreted as evidence that the Social Security "as implemented" also creates additional intergenerational risk by selectively providing certain cohorts with transfers that are too high or too low to be explained by economic outcomes of those cohorts. Simply stated, some cohorts were more "lucky" than others.

The selective treatment of certain cohorts was already suggested by Figures 1 and 2. Regressions in Table 2 provide a formal test; they divide cohorts into four groups[18] and estimate group-specific parameters β^W and β^S in the equation 1.[19] The relationship between differences in net transfers and wages/excess savings indeed varies widely between cohort groups; for wages, from positive and significant (cohort 1900-20) to negative, significant, and large in magnitude (cohort 1965-80); for excess savings, the disparities are even more pronounced.

3.2 Expected lifetime outcomes over time

While the preceding section estimated policy functions for outcomes observed at the end of the cohorts' lives, this section estimates analogous policy functions for outcomes that were expected at various points during the cohorts' lives. Expected lifetime outcome as of year s (e.g., the expected lifetime benefits) is defined as

[16] Standard errors are clustered by the "benchmark" cohort, and observations are weighted by the absolute value of the percentage difference between the cohorts' lifetime wages.

[17] The cohort fixed effects are also jointly statistically significant in both specifications.

[18] The groups are: 1900-20, 1925-40, 1945-60, 1965-85.

[19] The estimated equation is $\frac{T_i - T_k}{Y_k} = \sum_g \beta_g^W \lambda_g \frac{W_i - W_k}{Y_k} + \sum_g \beta_g^S \lambda_g \frac{S_i - S_k}{Y_k} + \lambda + \epsilon_{ik}$, where subscript g denotes the cohort group, β_g^W and β_g^S the group-specific coefficients, and λ_g are dummies denoting each cohort group.

$$B_i^s = B_i^{t \leq s} + E_s \left[B_i^{t>s} \right] \tag{2}$$

where $B_i^{t \leq s}$ denotes the present value of benefits received by cohort i until year s, and $E_s \left[B_i^{t>s} \right]$ denotes the present value of expected benefits received by cohort i after year s, the expectation being taken as of year s. The expected future contributions and benefits are computed under the assumption that the pension legislation valid in year s will remain valid through the cohort's lifetime.[20] Observing cohorts over time allows exploiting the panel structure of the data; for example, fixed cohorts effects remove any permanent differences in net transfers between cohorts that are due to permanent differences in their political influence. Assume that as of year t people had some expectations about their lifetime wages and savings, and the Social Security rules generated certain expected lifetime net transfer for each cohort. Economic shocks that were realized between years t and $t+s$ affected individual cohorts differently, say by increasing the expected lifetime income of cohort i relative to cohort k. Also as of year $t+s$, people form a new expectation about the lifetime net transfer, the difference from year t being a product of the default rules automatically adjusting the contributions and benefits to the shocks, and of possible reforms that may have been enacted between years t and $t+s$. If the Social Security system "as implemented" between years t and $t+s$ shares risk across generations, the net transfer to cohort i should fall relative to cohort k.[21]

Pension reforms are relatively infrequent and one would not expect economic shocks to be reflected in Social Security rules on a year-to-year basis. I therefore reduce the frequency to five-year intervals from 1940, 1945, etc to 2005. The expectation in each of these years used in the analysis is equal to a 5-year average of the current year and the preceding four years, e.g. the expectation for 1945 is the average over 1941-1945.

Figures 3-5 plot the changes in expected lifetime wages, excess savings, and net transfers over the five-year intervals for each cohort. The changes are expressed as fractions of expected lifetime income. The large negative change in 1945 for the older cohorts stems from the end of the rapid wage growth during the WWI, which, by construction of the adaptive expectations of future wages, fed into very high expected lifetime wages in 1940. There is visible heterogeneity between cohorts in the evolution of expected excess savings. As for the expected transfers, a distinct group of cohorts (1915-1935) experienced a series of positive changes during the 1970's

[20] The assumption need not imply that the parameters of the Social Security system will not change through the cohort's lifetime since the Social Security legislation sometimes stipulates changes in parameters in the future. Such changes are reflected in the expectation as of year s.

[21] Cohort k still may expect higher net transfer as of year $t+s$ than cohort j, either due to its higher political influence that carries through its lifetime or due to experiencing relatively worse economic shocks in the past.

and 1980's, while negative changes occur occasionally for the post 1950 cohorts and in later periods. The magnitude of changes in expected transfers is always below 5% of lifetime gross income.

Figure 6 shows how the relationship between expected lifetime net transfers and expected lifetime wages evolved over time. There are large differences in expected lifetime wages among cohorts until the 1960's but the expected lifetime net transfers are basically zero for all cohorts. In that sense, the Social Security "as implemented" resembled a defined contribution system, despite being formally a defined benefit. From 1970 onwards, ever larger disparities in expected lifetime net transfers begin to emerge, with younger cohorts being gradually made worse and worse off. However, such disparities do not generally reflect the differences in expected lifetime wages. During the 1980's and early 1990's there was in fact a "perverse" positive relationship between transfers and wages, as the recessions during that period reduced expected lifetime wages of the younger cohorts while net transfers to those cohorts were already lower than net transfers to older cohorts. Figure 7 depicts analogous relationships for expected lifetime excess savings. Since the 1970's there have been several years during which the lifetime net transfers were negatively related to excess savings.

The relationship between expected transfers, wages, and excess savings over time is quantified in regressions reported in Table 3, which are panel equivalents of equation 1. The unit of observation is the year of expectation s and a pairwise combination of cohorts i and k, with $i > k$. The kind of relationship being identified depends on the dummy variables included. Regression in column (1) includes fixed effects for "benchmark" cohorts k (which capture the average net transfer from cohort k vis-a-vis all younger cohorts across all years) and fixed effects for each year of expectation (which capture the average difference in net transfers for all younger cohorts i relative to cohort k in every year of expectation and reflect, for example, the gradually ameliorating generosity of the Social Security). The estimator is based on deviations from fixed effects both over cohorts and over years. The coefficient on wages is negligible; the coefficient on excess savings is negative (-0.076) and statistically significant at 1%. Column (2) contains fixed effects for each combination of the year of expectation and the "benchmark" cohort. It implicitly assumes a fixed "treatment" of a cohort k in a given year. The estimator is essentially a cross-sectional estimator of equation 1 but repeated many times. The results show again a negligible coefficient on wages and negative and significant (-0.150) coefficient on excess savings. These results are qualitatively similar to the results based on terminal outcomes but the estimated relationship between excess savings and net transfers is smaller by the order of magnitude.

Very different results emerge from column (3) which includes dummies for each combination of "comparison" and "benchmark" cohorts i and k. It implicitly assumes that each pair of cohorts has a fixed difference in net transfers "given" throughout their lives (reflecting, for example, the relative political power of the two cohorts), and then compares whether relative changes in wages or excess savings over time within this pair of cohorts yield to relative changes in net transfers. They do not; the coefficient on wages is again negligible while the coefficient on excess savings is actually positive (0.082) and significant. The results indicate that within a pair of cohorts, changes in transfers and wages actually go in the "wrong" direction.

A different perspective on risk sharing is obtained by estimating how aggregate shocks are translated into net transfers to particular cohorts. The change in aggregate lifetime wages $\overline{W}_t$ is defined as the change in the average expected lifetime wages, the average being taken across all cohorts that are alive and have not retired yet in the year of expectation t.[22] Equation 3 captures a question: If aggregate expected lifetime wages rise by 1 percent, does the "as implemented" Social Security translate this into a decrease or increase in expected lifetime net transfer to the young, the middle-aged, or the old? Specifically, I divide the cohorts in each year of expectation into four age groups: the non-working young (age 0-20), working young (25-45), working middle-aged (50-60), and the old who are also retired (65-80). $\lambda_a's$ then denote age group dummies and parameters β_a^W and β_a^S capture age-specific effects of the changes in aggregate expected lifetime wages and excess savings on expected lifetime net transfers of cohort i aged a.

$$T_{iat} = \sum_a \beta_a^W \lambda_a \overline{W}_t + \sum_a \beta_a^S \lambda_a \overline{S}_t + \lambda_t + \epsilon_{iat} \tag{3}$$

The estimates are presented in Table 4. The first specification includes year of expectation fixed effects while the other includes cohort fixed effects; the results are the same qualitatively but smaller in magnitude in the second specification. An increase in aggregate wages by 1 percent leads to a reduction in lifetime net transfer by 0.019 percent of lifetime income for the non-working young; an increase by 0.018 percent for the working young, and much larger increases (0.037 and 0.034 percent of lifetime income) by the working middle-aged and the old. The signs are exactly the opposite (and the absolute magnitudes very similar) for shocks to aggregate savings. In this regard, the Social Security "as implemented" redistributed aggregate shocks in a manner that was broadly consistent with optimal intergenerational risk sharing. Positive aggregate wage shocks benefit predominantly the young who will enjoy higher wages for a large fraction of their lives; hence the net transfers to the older generations rise. Positive shocks to returns to savings benefit predominantly the old who already have a large stock of savings; hence

[22] By construction of the data set, the expected lifetime wages or excess savings do not change after a cohort retires.

the net transfers to the older generations are reduced.

Regressions in Table 5 show in a different way the empirical difference between "stylized" and "as implemented" Social Security system. They estimate equation 3 but replace the age-group-specific effects with cohort-group-specific effects, the cohorts being divided into same groups as in the regressions with terminal outcomes. It thus measures whether aggregate shocks were being redistributed in a manner that selectively favored or disfavored certain cohorts defined by the year of birth, not by age. Aggregate shocks translated very differently into net transfers for cohorts born after 1945 than for the cohorts born earlier. The cohorts born after 1945 were relatively disfavored with respect to the aggregate shocks to wages but relatively favored with respect to the aggregate shocks to savings. These findings could again be rationalized by the fact that observations for these younger cohorts are based in large part on projections of the Social Security legislation valid in a given year into the future. Observations for the cohorts 1900-1945 are based in larger part on the outcomes already realized.

4 Conclusions

The pension literature has provided two rather contrasting views of risk in the pay-as-you-go pension systems: The normative theoretical literature on intergenerational risk sharing highlights the welfare improvements generated by the ability of the pay-as-you-go to transferring income across generation and gives valuable guidance on designing an optimal system. The literature on policy risk instead highlights that the pay-as-you-go systems are far from the ideal systems with stable rules. Legislative changes may and sometimes do reduce lifetime incomes of some people compared to what they were being promised by the preceding legislation.

This paper took a somewhat unusual, and basically positive, approach somewhere between the two literatures. It investigated whether the differences between cohorts in lifetime net transfers actually provided by the Social Security system can be rationalized by differences between cohorts in lifetime incomes, specifically in lifetime wages and returns on savings. The results give insights about the degree of intergenerational risk sharing that the Social Security system "as implemented" in the United States since the 1930's has actually provided. The most important finding, also for the future research, is that "as implemented" Social Security produces very different relationship between net transfers and economic outcomes than a "stylized" Social Security that is assumed to keep its parameters forever.

An important finding is that the net transfers from Social Security do reflect partially (and based on terminal outcomes, even fully) differences in the stock-market histories between cohorts. It is particularly striking given the fact that the Social Security rules do not contain explicit provisions linking stock market returns, contributions, and benefits. But Social Security appears to work "as if" it were designed to compensate the cohorts with lower lifetime returns on savings with higher lifetime net transfers. Changes in the generosity of the Social Security, as implemented through numerous reforms, appear to have taken to stock market histories into account in a manner that is broadly consistent with optimal risk sharing. On the other hand, net transfers do not reflect differences between cohorts in their lifetime wages, which is also surprising the fact that the Social Security index several variables to wage growth.

Some evidence shows that differences in net transfers across cohorts are to a large extent also driven by factors other than economic outcomes of the cohorts: statistically significant cohort-specific coefficients, cohort dummies or a perverse positive relationship between lifetime wages and net transfers for the oldest cohorts. Simply stated, some cohorts were more "lucky" than others. Should such unexplained differences indeed be unrelated to the cohorts' incomes, they can be interpreted as evidence that the policy risk embedded in the pay-as-you-go pension scheme also creates additional intergenerational risk.

References

[1] Auerbach, A. and R. Lee (2009). Welfare and generational equity in sustainable unfunded pension systems. NBER working paper 14682.

[2] Ball, Laurence, and N. Gregory Mankiw (2007): Intergenerational Risk Sharing in the Spirit of Arrow, Debreu, and Rawls, with Applications to Social Security Design. Journal of Political Economy, 115(3), 523-547.

[3] Bohn, Henning (2009). Intergenerational Risk Sharing and Fiscal Policy. Journal of Monetary Economics, 56(6), 805-816.

[4] Borgmann, C, and M. Heidler (2003). Demographics and volatile social security wealth: Political risk of benefit rule changes in Germany. CESifo working paper 1021.

[5] D'Amato M. and V. Galasso (2010). Political intergenerational risk sharing. Journal of Public Economics, 94(9-10), 628-637.

[6] Dusek, L. and J. Kopecsni (2008). Policy risk in action: pension reforms and social security wealth in Hungary, Czech Republic, and Slovakia. Czech Journal of Economics and Finance, 58(7-8), 329-357.

[7] Feldstein, M. and E. Ranguelova (2001). Individual Risk in an Investment-based social security system. American Economic Review, 91(4), 1116-1125.

[8] Gordon, H. and R. Varian (1988). Intergenerational risk sharing. Journal of Public Economics, 37(2), 185-202.

[9] Gottardi, P. and F. Kubler (2001). Social security and risk sharing. Journal of Economic Theory, 146(3), 1078-1106.

[10] Leimer, D.R. (1994). Cohort specific measures of lifetime net social security transfers. ORES working paper No. 59, Office of Research and Statistics, Social Security Administration.

[11] Leimer, D.R. (2007): Cohort Specific Measures of Lifetime Social Security Transfers. ORES working paper No. 110, Office of Research, Evaluation, and Statistics, Social Security Administration.

[12] McHale, John (2001): The Risk of Social Security Benefit-Rule Changes: Some International Evidence. In Campbell, John Y. and Martin Feldstein (eds.), Risk Aspects of Investment-Based Social Security Reform. University of Chicago Press.

[13] Poterba, James, Steven Venti, and David Wise (1998). Implications of Rising Personal Retirement Saving. In Wise, D. (ed.), Frontiers in the Economics of Aging, University of Chicago Press.

[14] Rengel, A. and R. Zeckhauser (2001). Can market and voting institutions generate optimal intergenerational risk sharing? In Campbell, John Y. and Martin Feldstein (eds.), Risk Aspects of Investment-Based Social Security Reform, University of Chicago Press.

[15] Shoven, John and Sita N. Slavov (2006). Political Risk versus Market Risk in Social Security. NBER Working Paper No. 12135.

[16] Wagener, A (2004). On intergenerational risk sharing within social security schemes. European Journal of Political Economy, 20(1), 181-206.

Figure 1: Lifetime wages

Figure 2: Lifetime excess savings

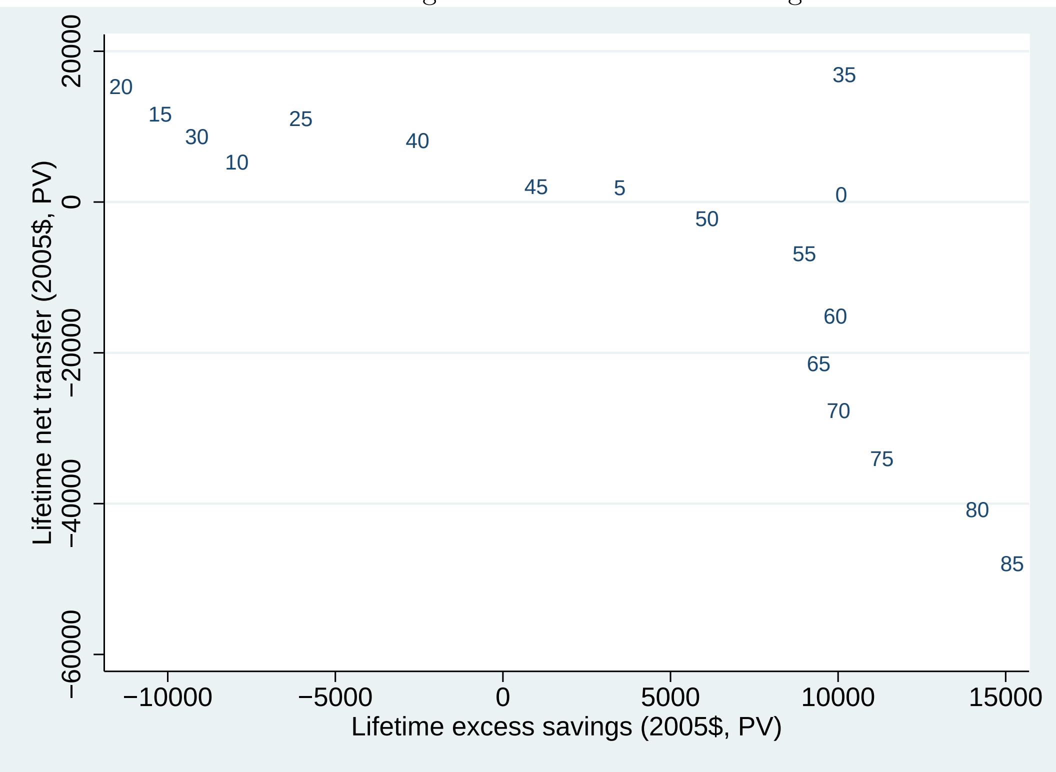

Figure 3: Changes in expected lifetime wages over time

Figure 4: Changes in expected lifetime excess savings over time

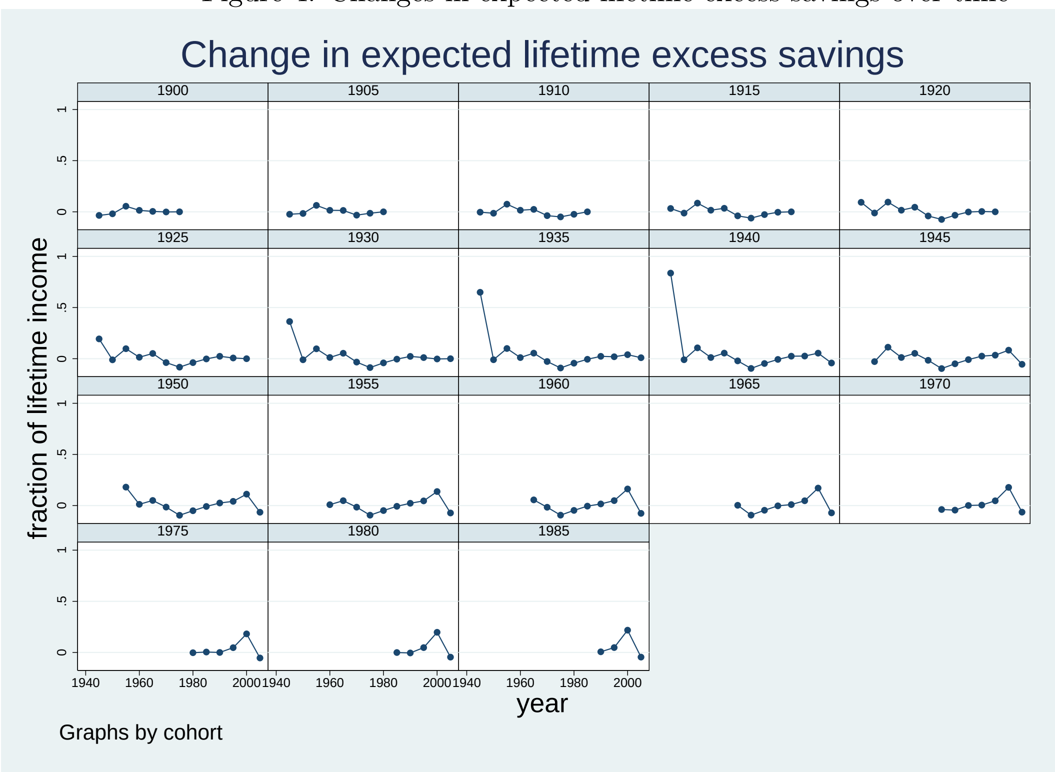

Figure 5: Changes in expected lifetime net transfer over time

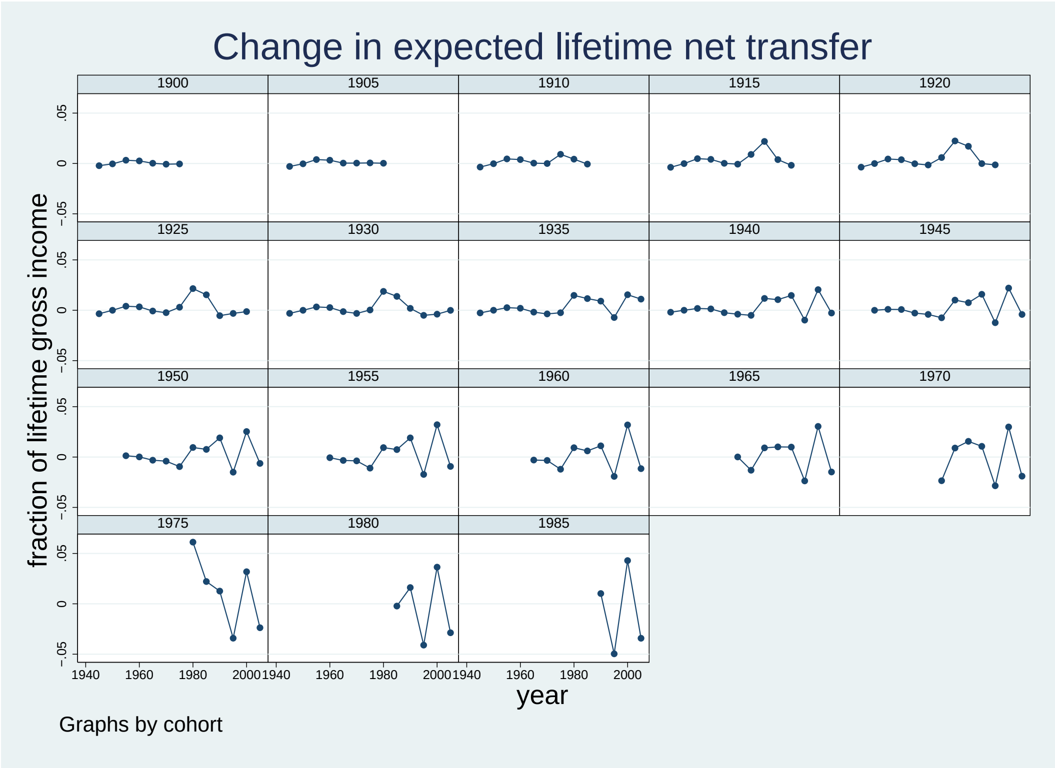

Figure 6: Expected lifetime net transfers and wages over time

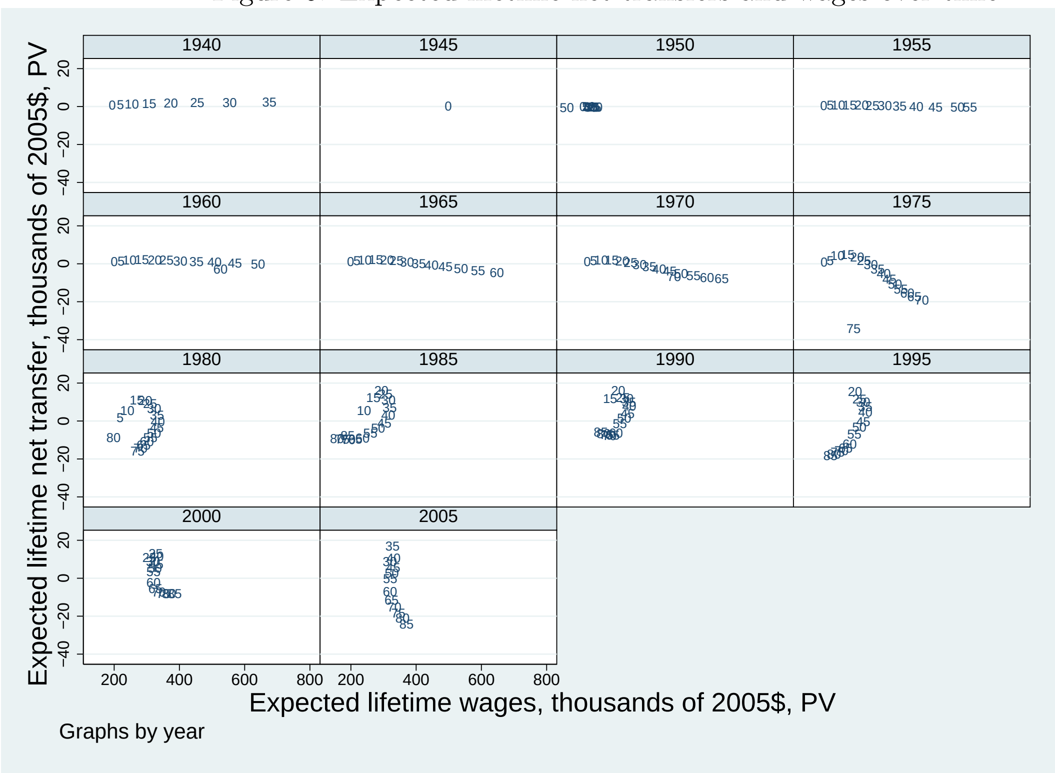

Figure 7: Expected lifetime net transfers and excess savings over time

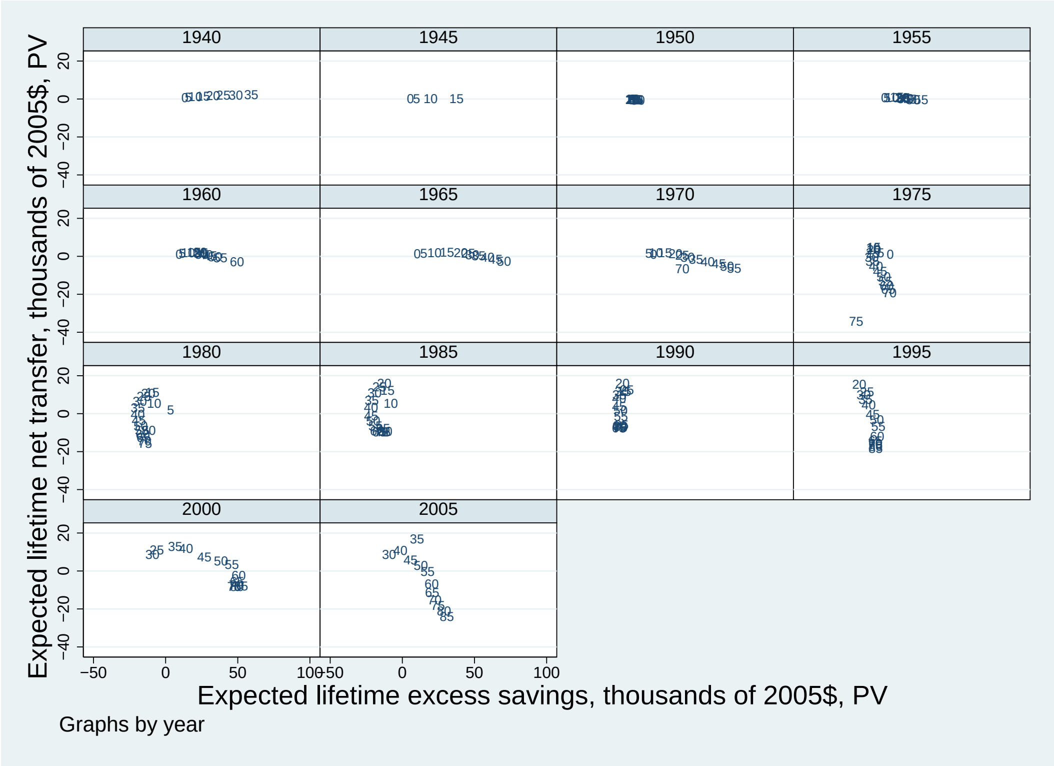

Table 1: Terminal outcomes

difference in lifetime net transfers	(1)	(2)	(3)
difference in lifetime wages	0.138	-0.067	-0.118
	(0.032)**	(0.032)*	(0.036)**
difference in lifetime excess savings	-2.027	-0.602	-1.307
	(0.037)**	(0.192)**	(0.117)**
constant	-0.015	-0.014	0.002
	(0.000)**	(0.010)	(0.013)
benchmark cohort dummies	yes	no	no
comparison cohort dummies	no	yes	no
observations	153	153	153
R-squared	0.610	0.961	0.453

robust standard errors in parentheses.

all differences normalized by the benchmark cohort's lifetime income.

* significant at 5%; ** significant at 1%

unit of observation: pairwise combination of cohorts $i - k$, where $i > k$,

terminal outcomes

Table 2: Terminal outcomes - cohort group-specific effects

difference in lifetime net transfers		(1)	(2)
	cohort 1900-20	-0.158	0.138
		(0.033)**	(0.033)**
difference in lifetime wages	cohort 1925-40	0.593	-0.225
		(0.375)	(0.202)
	cohort 1945-60	0.833	-0.374
		(0.174)**	(0.112)**
	cohort 1965-85	-1.789	-2.098
		(0.745)*	(0.450)**
	cohort 1900-20	-1.335	-2.020
		(0.102)**	(0.036)**
difference in lifetime excess savings	cohort 1925-40	-1.117	-0.457
		(0.372)**	(0.196)*
	cohort 1940-60	-3.799	-2.435
		(0.877)**	(0.636)**
	cohort 1965-85	0.614	2.395
		(0.877)	(0.767)**
cohort dummies		no	yes
observations		153	152
R-squared		0.47	0.61

robust standard errors in parentheses.

all differences normalized by the benchmark cohort's lifetime income.

* significant at 5%; ** significant at 1%

unit of observation: pairwise combination of cohorts $i - k$, where $i > k$,

Table 3: Expected outcomes over time

difference in lifetime net transfers	(1)	(2)	(3)
difference in lifetime wages	0.001	0.001	-0.001
	(0.000)**	(0.000)**	(0.000)**
difference in lifetime excess savings	-0.076	-0.150	0.082
	(0.018)**	(0.028)**	(0.012)**
constant	0.013	0.022	-0.148
	(0.002)**	(0.003)**	(0.001)**
comparison cohort dummies	yes	no	no
comparison $\times$ benchmark cohort dummies	no	no	yes
year of expectation dummies	yes	no	no
cohort $\times$year of dummies	no	yes	no
observations	1176	1026	1026
R-squared	0.47	0.55	0.73

robust standard errors in parentheses

all differences normalized by the benchmark cohort's lifetime income

* significant at 5%; ** significant at 1%

unit of observation: pairwise combination of cohorts $i-k$, where $i>k$, and year of expectation

Table 4: Sharing of aggregate shocks by age groups

lifetime transfer as a share of lifetime income		(1)	(2)
	age 0-20	-0.019**	-0.011**
log of aggregate	age 25-45	0.018**	0.001
wages	age 50-60	0.037**	0.019**
	age 65-80	0.034**	0.019**
	age 0-20	0.027**	0.013**
log of aggregate	age 25-45	-0.020*	-0.002
savings	age 50-60	-0.041**	-0.021**
	age 65-80	-0.035**	-0.021**
year of expectation dummies		yes	no
cohort dummies		no	yes
observations		186	186
R-squared		0.73	0.81

** significant at 1%, * significant at 5%

Table 5: Sharing of aggreagate shocks by cohort groups

lifetime transfer		(1)	(2)
	cohort 1900-20	-0.002	0.012*
log of aggregate	cohort 1925-40	-0.008	-0.010
wages	cohort 1945-60	-0.029**	-0.029**
	cohort 1965-85	-0.033**	-0.013
	cohort 1900-20	-0.008	-0.017*
log of aggregate	cohort 1925-40	0.009	-0.009
savings	cohort 1945-60	0.033**	-0.032**
	cohort 1965-85	0.034**	0.029
year of expectation dummies		yes	no
cohort dummies		no	yes
observations		186	186
R-squared		0.79	0.73

** significant at 1%, * significant at 5%