

This PDF is a selection from a published volume from the
National Bureau of Economic Research

Volume Title: The Distributional Aspects of Social Security
and Social Security Reform

Volume Author/Editor: Martin Feldstein and Jeffrey B. Liebman,
editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-24106-8

Volume URL: <http://www.nber.org/books/feld02-1>

Conference Date: October 21-23, 1999

Publication Date: January 2002

Title: The Impact of Social Security and Other Factors on
the Distribution of Wealth

Author: Jagadeesh Gokhale, Laurence J. Kotlikoff

URL: <http://www.nber.org/chapters/c9749>

The Impact of Social Security and Other Factors on the Distribution of Wealth

Jagadeesh Gokhale and Laurence J. Kotlikoff

3.1 Introduction

As documented in Feldstein (1976) and Auerbach et al. (1995), the postwar period has witnessed a dramatic increase in the annuitization of the resources of America's elderly. Indeed, between 1960 and 1990, the annuitized share of resources of older men doubled, and that of older women quadrupled. Social Security is the main force behind this process, but Medicare and Medicaid as well as private pensions play an important role in replacing household wealth accumulation with survival-dependent old-age resource streams. Gokhale, Kotlikoff, and Sabelhaus (1996) suggest that the increased annuitization may explain the significant postwar rise in the propensity of the elderly to consume their remaining lifetime resources. Gokhale et al. (2001) consider a related point, namely, that increased annuitization will reduce bequests. This is particularly the case for lower-income households, who face Social Security taxation on all of their earnings and for whom Social Security replaces a higher share of preretirement income. Thanks to Social Security, these households have less to save, and less reason to save, than upper-income households. Thus, they arrive at old age with relatively little net worth available to bequeath. In differentially disenfranchising the children of the poor from the receipt of

Jagadeesh Gokhale is a senior economic advisor with the Federal Reserve Bank of Cleveland. Laurence J. Kotlikoff is professor of economics at Boston University and a research associate of the National Bureau of Economic Research.

The opinions expressed here are those of the authors and not necessarily those of Boston University, the Federal Reserve Bank of Cleveland, or the National Bureau of Economic Research. Laurence J. Kotlikoff thanks the National Bureau of Economic Research, the National Institute of Aging, and Boston University for research support.

inheritances, Social Security may be materially altering the distribution of wealth.

This chapter builds on Gokhale et al. (2001) in trying to understand how Social Security and other factors affect wealth inequality. Specifically, this chapter uses lifetime earnings data from the Panel Study of Income Dynamics (PSID) to calibrate Gokhale et al.'s (2001) bequest simulation model. The model features random death based on age-sex mortality probabilities, random fertility, assortative mating, heterogeneous human capital endowments, heterogeneous rates of return, the inheritability of human capital, progressive income taxation, and the partial annuitization through Social Security of households' retirement savings. Agents have no bequest motive. They live for a maximum of eighty-eight years, the first twenty-two as children, the second twenty-two as young married adults who have children, the third twenty-two as middle-aged adults raising children, and the last twenty-two as elderly individuals who die at random. Surviving spouses inherit all their partners' wealth but bequeath their own wealth in equal portions to their children.

Gokhale et al. (2001) use synthetic data generated by CORSIM—a micro simulation model—in conjunction with cross-section earnings data from the Survey of Consumer Finances (SCF) to calibrate the distribution of lifetime labor earnings. Surprisingly, their simulated distribution of wealth held by those reaching retirement closely matches that observed in the SCF. This is particularly true at the upper tail of the wealth distribution. In addition to matching the wealth distribution, Gokhale et al. decompose the influence of various factors, including Social Security, on wealth inequality. They show that, in the absence of Social Security, inheritances would reduce slightly the degree of wealth inequality. In Social Security's presence, however, they find that inheritances significantly exacerbate wealth inequality.

While intriguing, Gokhale et al.'s results should be viewed with caution, given the synthetic generation of lifetime earnings upon which they rely. This chapter attempts to alleviate that shortcoming by considering actual lifetime earnings as reported in the long PSID panel. The availability of actual lifetime earnings also permits a more accurate calibration of the degree of assortative mating and inheritability of skills. These advantages must be set against a key disadvantage: The PSID does not adequately capture the skewness in the upper tail of the distribution of lifetime earnings. Since the upper tail of the distribution of lifetime earnings plays such a key role in determining the upper tail of the distribution of wealth, this paper focuses on the shape of the distribution of wealth in the bottom half of the wealth distribution.

Because we are using the PSID to generate the distribution of lifetime earnings, and because the PSID data suggest a rather minor degree of assortative mating, our simulated distribution of wealth holdings among

households reaching retirement is much less skewed than that reported in Gokhale et al. (2001). It is also less skewed than the 1995 SCF and 1984 PSID wealth distributions of married households aged sixty to sixty-nine.

Although we are less successful than Gokhale et al. (2001) in matching the upper tail of the wealth distribution, we continue to confirm Feldstein's (1976) finding that Social Security plays an important role in making wealth holdings less equal. The Gini coefficient for our simulated distribution of wealth is 16 percent higher in the presence of Social Security than in its absence, and Social Security raises the share of wealth held by the wealthiest 10 percent of retiring households by almost one quarter.

The paper proceeds in section 3.2 with a brief literature review. Section 3.3 draws heavily on Gokhale et al. (2001) in describing the model and proceeds to discuss the model's calibration. Section 3.4 presents results, and section 3.5 draws conclusions.

3.2 Related Literature

Harbury and Hitchens (1979) and Menchik (1979) are well-cited studies that connect the wealth of sons to the estates of their fathers. While suggestive, their data do not necessarily imply that inherited wealth is a source of inequality.¹ Nevertheless, they reinforced Meade's (1976) view that inheritances exacerbate wealth inequality—a conclusion supported by the theoretical model of Wilhelm (1997). Stiglitz (1969) and Atkinson and Stiglitz (1980) offered an alternative perspective: They showed that, under certain assumptions, a stable, egalitarian distribution of wealth would emerge if inheritances were distributed evenly among all of one's children.

The models by Becker and Tomes (1979) and Tomes (1981) go beyond these studies in examining the joint role that inheritances of financial wealth and earning power (human capital) play in determining whether intergenerational transfers are equalizing. Laitner (1979a, b) constructs a utility-maximizing framework in which parents care about both their own and their children's consumption, bequests must be nonnegative, and there is no inheritability of human capital. He shows that an equilibrium wealth distribution exists and that inheritances are equalizing if there is no assortative mating—an issue first examined by Blinder (1973). The studies of Meade (1964), Stiglitz (1969), Pryor (1973), Atkinson and Harrison (1978), and Atkinson (1980) also stress the role of imperfect correlation of spouses' inheritances in equalizing the distribution of inheritances.²

1. Their data record the total value of the father's estate, rather than the amount actually inherited by the son. Their findings may be explained, in large part, by the sons' inheritances of their fathers' human capital.

2. Theoretical work on taxation provides additional grounds for believing that inheritances are equalizing. Becker and Tomes (1979) and Atkinson and Stiglitz (1980) show that inheritance taxation can increase income inequality. However, if there are incomplete markets,

Early simulation studies of wealth inequality include Atkinson (1971), Flemming (1976), Oulton (1976), and Wolfson (1977). Their central finding is that, absent bequests, the life-cycle model is unable to explain the upper tail of the wealth distribution. Wolfson (1979) and Davies (1982) add bequests, specifically desired bequests, and generate much more realistic skewness in the distribution of wealth.

Flemming (1979) is the closest antecedent to our work. He, too, considers earnings heterogeneity and the inheritability of skills, but not marriage, assortative mating, or heterogeneity in the number and spacing of children. Flemming finds that wealth is much more unequally divided than earnings and that both unintended and intended bequests can markedly increase wealth inequality. Although many of our findings agree with Flemming's, we find that, in the absence of Social Security, unintended bequests serve to slightly reduce intragenerational wealth inequality.

3.2.1 Empirical Support for Excluding Altruism

Although Laitner and Juster (1996) provide some limited support for altruism, most recent studies do not. To be precise, Boskin and Kotlikoff (1985), Altonji, Hayashi and Kotlikoff (1992, 1997), Abel and Kotlikoff (1994), Hayashi, Altonji, and Kotlikoff (1996), Gokhale, Kotlikoff, and Sabelhaus (1996), Wilhelm (1996), and Hurd (1992) show that (a) the distribution of consumption across cohorts is very strongly dependent on the cross-cohort distribution of resources, (b) the distribution of consumption within extended families is very strongly dependent on the distribution of resources within extended families, (c) that taking a dollar from a child and handing it to parents who are actively transferring income to that child leads the parent to hand back only thirteen cents to the child, (d) that the very major postwar increase in the annuitization of the resources of the elderly has not been even partially offset by an increase in their holdings of life insurance, (e) that the vast majority of bequests are distributed equally among children independent of their economic needs, and (f) the presence of children does not influence postretirement dissaving. Individually and as a group, these studies constitute very strong evidence against intergenerational altruism, suggesting that most bequests may be unintended or motivated by nonaltruistic considerations—the modeling assumption made here.³

such as the market for educational loans considered by Loury (1981), redistributive taxation of bequests can reduce intragenerational inequality.

3. Of course, the Bill Gateses and Warren Buffets of the world are not making unintended bequests. Unlike most people, they may be able to afford their altruism because of their high levels of income. Stated differently, many households may be altruistic, but their degrees of altruism are so small that, given their resources, they choose not to allocate resources toward intentional, altruistically motivated bequests.

It is frequently observed that retired consumers save rather than dissave. *Prima facie* this might indicate a bequest motive in some form (Hurd 1990). However, Gokhale, Kotlikoff, and Sabelhaus (1996) and Miles (1997) point out that when wealth is calculated to include the capitalized value of Social Security receipts, it falls throughout retirement. Hurd also concludes from a careful analysis of panel and cross-section data that the evidence on wealth changes is consistent with the life-cycle hypothesis and the view that bequests are accidental.

Meade (1966) and Flemming (1976) provide a final reason to doubt the prevalence of altruistically motivated bequests. They pointed out that anything less than very strong altruism would not suffice to generate ubiquitous and significant bequests because the lifetime incomes of children significantly exceed those of their parents.

3.3 The Model and Its Calibration

This section describes the model's demographic structure, skill allocation, determination of inheritance, and consumption and saving behavior.

3.3.1 Demographics

Agents can live for eighty-eight periods, the first twenty-two of which they spend as children whose consumption is financed by their parents. All events, including earnings, consumption, marriages, births, deaths, and wealth transfers, occur at the end of each period. Agents marry on their twenty-second birthday. They give birth to children at ages twenty-two through forty-three, depending on their draw from a "birth matrix" described below. They also enter the work force on their twenty-second birthday and work through age sixty-six. They face positive probabilities of dying between ages sixty-seven and eighty-eight. The probability of an agent's dying on her eighty-eighth birthday, given that she has lived to that date, is 100 percent. The probabilities of dying at ages sixty-seven through eighty-seven are taken from U.S. mortality statistics.

The number, sexes, and timing of children born to each couple are determined randomly. This distribution is aligned to ensure that 2,000 males and 2,000 females are born each year. There is no population growth, so total annual births remain fixed through time.

3.3.2 Consumption and Saving Behavior

Agents' expected utilities are time-separable isoelastic functions of their own current and future consumption as well as that of their children through age twenty-two. Consider, as an example, the expected utility (EU) of a couple, aged twenty-three, who will have two children, one when the couple are aged twenty-five and the other when they are twenty-eight:

$$(1) \quad \text{EU} = \sum_{a=22}^{a=87} \beta^{a-22} (p_{ha} c_{ha}^{1-1/\sigma} + p_{wa} c_{wa}^{1-1/\sigma}) \\ + \delta \sum_{a=25}^{a=46} \beta^{a-22} c_{k1a}^{1-1/\sigma} + \delta \sum_{a=28}^{a=49} \beta^{a-22} c_{k2a}^{1-1/\sigma} .$$

In equation (1), the first summation considers the utility of each spouse from his or her own consumption at each possible age to which each could live. The second two summations consider the utility that the couple derive from the consumption of their two children. The terms c_{ha} , c_{wa} , c_{k1a} , and c_{k2a} refer, respectively, to the consumption of the husband, wife, first child, and second child when the couple is age a . The term β is the time-preference factor, σ is the intertemporal elasticity of substitution, and δ is the child consumption weighting factor.

As σ approaches zero, households become more and more reluctant to consume smaller amounts in the future than they consume in the present. Since the inverse of σ is the household's coefficient of relative risk aversion, a value of close to zero translates into a coefficient of risk aversion close to infinity. In our simulations, we assume that σ is very close to zero. Hall (1988) reports that there is "no strong evidence that the elasticity of intertemporal substitution is positive. Earlier findings of substantial positive elasticities are reversed when appropriate estimation methods are used."

Although Hall's findings are based on a highly stylized model, assuming that σ is very close to zero enormously simplifies the consumption decisions of our model's households. First, this assumption in conjunction with the assumption of a time preference rate equal to the interest rate means that households seek to maintain the same level of consumption over time for each spouse. Households also seek to maintain a constant level of consumption for their children, while they are children. Given the value of δ , this child consumption level equals 40 percent of the parental consumption level.

Most important, however, our assumption that σ is very close to zero means that households only consider their safe resources in deciding how much to consume at each point in time. Thus, households who expect to receive an inheritance, but are uncertain of receiving one (because all of their parents may live to age eighty-eight), will ignore this potential source of future income in making their current consumption and saving decisions.

At each point in time, married households will calculate the number of years of remaining life, multiply this amount by two (to take into account the presence of both spouses) and then add to the resulting value 0.4 times the number of years of consumption of their children. This total number of effective adult consumption years is then divided into the household's safe resources to determine consumption per effective adult. The household's safe resources consist of its wealth (which may reflect the receipt of

past inheritances) plus the present value of its remaining lifetime labor earnings, which, by necessity, is assumed to be certain. Given the level of consumption per effective adult, it is straightforward to calculate total household consumption and subtract it from total household income to determine household saving.

We want to emphasize that inheritances affect consumption behavior, but only once they are received. There is no consumption out of potential future inheritances. Instead, households consider the worst-case scenario at each point in time and formulate their consumption and saving plans accordingly. Were we to assume a positive value of σ , households would take a gamble and consume more in the present in anticipation of a possible future inheritance. However, their decision as to how much to consume would be extraordinarily complex. The reason is that at certain ages they would have to take into account not simply their own resources, including their own wealth, but also those of their parents and grandparents, assuming their grandparents are still alive. Take, for example, a twenty-five-year-old couple with two sets of living parents and four sets of living grandparents. In deciding how much to consume, the household has to consider its own current wealth level as well as the wealth levels of all six parental and grandparental households. Formally, the dynamic program that the household must solve to determine how much to consume involves up to seven state variables, namely all seven of these wealth levels.⁴ Unfortunately, solving dynamic programs with seven state variables appears to be beyond the capacity of current computers.⁵

3.3.3 Assortative Mating

Agents are assigned their marriage partners at age twenty-two on a partly systematic basis in which the probability of a match is higher the closer are the respective skill ranks of the two partners. Our assortative mating works as follows. We first construct for each year of the simulation a vector \mathbf{A} containing 2,000 random numbers drawn from a uniform distribution with support $(0, 1)$. Vector \mathbf{A} is sorted in ascending order. Second, a vector \mathbf{B} is constructed containing 2,000 random numbers from the same distribution but is left unsorted. Third, vector $\mathbf{C} = \alpha\mathbf{A} + (1 - \alpha)\mathbf{B}$ is constructed to yield a new vector of numbers between zero and one. Fourth, we associate with the first element of the \mathbf{C} vector the integer 1, with the second element the integer 2, and so on up to the last element of \mathbf{C} , to which we associate the integer 2,000. Fifth, we sort \mathbf{C} such that \mathbf{C} ends up in ascending

4. We say “up to” because during years in which the household is aged sixty-six and over, it has neither living parents nor living grandparents, and during years in which the household is aged forty-four through sixty-five, it has no living grandparents.

5. The fact that even supercomputers would have difficulty solving this problem in a reasonable amount of time raises the question of how mere mortals can actually deal with this complexity.

order and denote the associated integers as vector \mathbf{R} . This vector is used to assign the husband's skill rank, j , for each female skill rank, where the female skills are aligned in ascending order), $i = 1 \dots 2,000$.

For example, suppose that, for a given year, after the \mathbf{C} vector has been generated and sorted in ascending order, the first three integers associated with the first three elements of the \mathbf{C} vector are 3, 51, and 1,290. Then the most skilled twenty-two-year-old female in that year will be married to the 3rd most skilled twenty-two-year-old male, the 2nd most skilled twenty-two-year-old female will be married to the 51st most skilled twenty-two-year-old male, and the 3rd most skilled twenty-two-year-old female will be married to the 1,290th most skilled twenty-two-year-old male. The parameter α is chosen to reproduce the rank correlation coefficient between husbands' and wives' lifetime earnings, which we estimate in the PSID.

3.3.4 Inheritance of Skills

The procedure just described for correlating husbands' and wives' lifetime earnings is also used to determine the correlation between fathers' and sons' lifetime earnings as well as the correlation between mothers' and daughters' lifetime earnings. In each case, the relevant parameter is calibrated to reproduce the rank correlation coefficient observed between sons' (daughters') and fathers' (mothers') lifetime earnings. We use the assumption that if sons (daughters) inherit skills, they do so from their fathers (mothers), which, while both politically and genetically incorrect, maintains tractability of the distribution of earnings through time of males and females.

3.3.5 Fertility

An initial population (at time $t = 0$) of 4,000 thousand individuals (2,000 males and 2,000 females) was created for each age between zero and eighty-seven. This was done as follows: First, a matrix of "birth ages" was derived from a fertility simulation of CORSIM—a dynamic microsimulation model of the U.S. economy described in Caldwell et al. (1999). The simulation considered 40,434 females born between 1945 and 2000 and recorded their ages of giving birth if those ages fell between ages twenty-two and forty-three. For each female in our CORSIM sample, we stored this information in our CORSIM birth matrix, which accommodates a maximum of ten birth ages, five for male and five for female births. Thus, the matrix has 40,434 rows and ten columns. Table 3.1 shows the distribution of females in the CORSIM matrix by number and sex of births.

Because computer memory limitations allowed us to process only 4,000 individuals in each year of birth, we needed to pare down our birth matrix in such a way as to end up with a modified birth matrix that contains exactly 2,000 male births and 2,000 female births. We started by selecting

Table 3.1 **Distribution of Females by Number and Sex of Births at Ages 22 through 43**

	Female Births					
	0	1	2	3	4	5
Male Births	<i>CORSIM Data</i>					
0	20.68	12.94	6.96	1.49	0.33	0.10
1	13.64	14.77	5.32	1.38	0.35	0.13
2	7.83	5.06	2.16	0.70	0.24	0.12
3	1.87	1.38	0.64	0.23	0.12	0.05
4	0.41	0.33	0.17	0.12	0.08	0.00
5	0.12	0.13	0.11	0.05	0.00	0.00
	<i>Birth Matrix Used in Simulation</i>					
0	0.00	13.00	9.00	2.20	0.55	0.15
1	16.05	22.40	6.45	1.60	0.50	0.00
2	12.20	7.15	2.50	0.80	0.00	0.00
3	2.85	0.95	0.75	0.00	0.00	0.00
4	0.40	0.40	0.00	0.00	0.00	0.00
5	0.10	0.00	0.00	0.00	0.00	0.00

Source: Authors' calculations based on CORSIM birth matrix.

Notes: CORSIM data represent females with at least one but fewer than five births numbered 40434. Selections from the birth matrix are restricted to females with one through five births, male and female (shown in boldface type).

2,000 rows from the birth matrix. We performed selection at random and without replacement, except that rows containing more than five births were excluded. The total number of births in the selected 2,000-row matrix exceeded 4,000. Hence, we randomly eliminated male and female births in the rows of this matrix for rows containing more than one birth until we were left with precisely 2,000 male and 2,000 female births. This guaranteed that the 2,000 rows of the final birth matrix would generate exactly 2,000 female and 2,000 male births. Table 3.1 shows the distribution of females by the number and sex of their births in the birth matrix used in the simulation.

3.3.6 Populating the Model at Time Zero

We populated our model by first creating 2,000 male and 2,000 female old adults for each age between sixty-seven and eighty-eight. These males and females were then married to each other sequentially. Some of these oldsters were treated as dead when we initiated the simulation, but we needed to include their ghosts at this stage of our process of populating the model in order to establish complete family trees. Marriage was allowed only between people of the same age to be consistent with our assumption that marriage occurs at age twenty-two (i.e., that initial oldster

males married initial oldster females when they were twenty-two and their brides were twenty-two). Family relationships were established by exchanging identification numbers.

In our next step, we drew from the 2,000 rows of the birth matrix at random (but without replacement) the middle-aged and young-adult children of the initial oldsters, ranging in age from twenty-four through sixty-six. In this process, we do not permit oldsters to bear children in their twilight years; rather, we are retrospectively considering the births of the initial oldsters when they were in their childbearing years.

Given that females give birth between the ages of twenty-two and forty-three, oldsters aged eighty-eight at the initiation of our simulation ($t = 0$) have children who are aged forty-five through sixty-six; oldsters aged eighty-seven at $t = 0$ have children aged forty-four through age sixty-five; and so on, until we reach oldsters aged sixty-seven at $t = 0$ who would have children aged between twenty-four and forty-five. Thus, at this stage of our populating procedure, exactly 4,000 (the full complement of) forty-five-year-olds and less than 4,000 thousand individuals at other ages between twenty-four and sixty-six have been created. The reason is that everyone (including oldster ghosts) who could have given birth to forty-five-year-olds has been considered, but not everyone who gave birth to those between ages twenty-four and forty-four and those between ages forty-six and sixty-six has been considered.

Because at this stage there are fewer than 4,000 middle-aged males and females at ages forty-six to sixty-six, additional middle-aged males and females are created for a total of 4,000 for each of these age groups. Next, all middle-aged males and females (those aged forty-five through sixty-six) were married at random, making sure that siblings were not married to each other. Next, the children of middle-aged adults were created, again taking draws without replacement from the birth matrix for females of a given age and then doing the same for females of another age until all females aged forty-five through sixty-six had been considered. The children produced by this process range in age from two through forty-four.⁶ Given that we have already created the children of the $t = 0$ oldsters, the addition of these children leave us with exactly 2,000 males and 2,000 females aged twenty-three through forty-four—the young adults. The procedure just described was also used to marry the young adults.

The next step in the creation of the initial population was creating the children of the $t = 0$ young adults that were born at $t = -1$ or earlier. Each young-adult female was assigned a row of the birth matrix at random

6. Sixty-six-year-olds have children aged between twenty-three and forty-four; sixty-five-year-olds have children aged between twenty-two and forty-three; and so on through forty-five-year-olds, who have children aged between two and twenty-three.

without replacement, and children were created for all birth ages less than the age at $t = 0$ of the female in question. For example, a forty-four-year-old female's children were created for birth ages between twenty-two and forty-three, but a twenty-three-year-old's children are created only if her birth row assignment contains a birth age of twenty-two. That is, children that will be born at $t = 0$ or later were not created as yet. At the end of this process, exactly 2,000 males and 2,000 females had been created for each age between one and eighty-eight. The final step in creating the initial population was to kill off oldsters (i.e., make the ghosts disappear) according to their cumulative mortality probabilities.⁷

3.3.7 Populating the Model Through Time

In populating the model through time we do the following. First, for $t = 0$, we allocate at random and without replacement a row from the birth matrix to all twenty-two-year-old females. Second, we marry twenty-two-year-old males and females at random, or according to the assortative mating procedure described above. Third, we have females aged twenty-two to forty-three give birth as determined by their assigned birth matrix row, creating 2,000 newborn (zero-year-old) males and 2,000 newborn (zero-year-old) females. Fourth, we kill off oldsters at random according to the conditional probability of dying at their respective ages. The wealth of these oldsters is transferred to the surviving spouse or children. Finally, we age everyone by one year.

3.3.8 Using the Panel Study of Income Dynamics to Calculate Lifetime Earnings

The PSID began in 1968 with a representative sample of 5,000 U.S. households. The PSID has re-interviewed (or attempted to re-interview) the individuals from those households every year since that time, regardless of their demographic status, living arrangements, etc. Children of original PSID respondents have been followed after they have left their parents' households.

Forming longitudinal labor income profiles involves the following steps: First, labor income, family number, and sex variables for "head" and "wife" are extracted from the family file for year x and merged with the individual file by year- x family number, retaining only those observations that appeared in year x .⁸ Step 1 is repeated separately for each year from

7. The mortality probabilities are based on U.S. mortality tables. Conditional mortality probabilities below age sixty-seven are set to zero, and the conditional mortality probability at age eighty-eight is set to unity. The probability of dying at age a , d_a , is calculated as $d_a = (1 - \sigma_a) \prod_{s=67}^{a-1} \sigma_s$, where σ_s is the conditional probability of surviving at age s .

8. The PSID's file structure comprises one cross-year individual file and several single-year family files. The cross-year individual file contains records for all individuals that ever ap-

1968 to 1993 to create twenty-six separate data files. Each data file is sorted by “family interview number” and “person number” respectively, and all are merged together to form a single “cross-year individual” file containing individual longitudinal labor income profiles.

The PSID reports annual labor income in nominal dollars. To place the income earned in different years on a comparable basis, we divide each year’s values by the ratio of that year’s Social Security Administration’s average wage index to the average wage index in 1997. Using this wage index to rescale nominal labor income adjusts not only for inflation but also for labor productivity growth. We have selected 1997 as base year because that year’s federal income tax rates and brackets are used to implement progressive income taxation in our simulation. Our motivation for controlling for productivity growth in forming lifetime earnings is simple: Our model does not account for productivity growth. Our next step is to sort the cross-year individual file by sex.⁹

Not all individuals appear in the PSID in each year between 1968 and 1993. Moreover, some individuals appear as “head” or “wife” for just a few years. To ensure a minimum number of data points from which to form lifetime earnings, we retained in our sample only those individuals who appear for at least ten years as “head” or “wife” between the ages of twenty-three and sixty-six.

Since the PSID spans only twenty-six years, it cannot provide us with earnings profiles for the full complement of forty-four years (from age twenty-three to age sixty-six) of earnings—the lifetime earnings profile as required by our simulation. Hence, we adopt a procedure for extrapolating each individual’s labor income both backward and forward as required. First, we calculate average labor incomes for all males who appeared in the sample (as “head” or “wife”) at each age between twenty-three and sixty-six.¹⁰ This provides us with a benchmark age-earnings profile for males—M23–M66. This procedure is repeated for females to generate a corresponding benchmark profile for females—F23–F66.

Next, for each individual we calculate two average labor incomes for each person. A1 is the average labor income over the first five years of the

peared in the survey—whether these individuals were members of the original sample of households or households that formed after 1968. This individual file contains a “person number” variable indicating the response status for each individual, which can be associated with the family number to create a unique person identifier. Unfortunately, the cross-year individual file does not contain labor income for “head” and “wife” for all of the years. Hence it was necessary to select these variables from the single-year family files and merge them together to form individual longitudinal labor income profiles.

9. We classify appropriately female “heads” and male “wives.”

10. Thus, those who appeared in the sample at their age twenty-three are included in the calculation of the average labor income for age twenty-three; those who appeared in the sample at age twenty-four are included in the calculation of the average labor income at age twenty-four, and so on through age sixty-six.

person's presence in the sample, and A_2 is the average labor income over the last five years that the person is present in the sample (as "head" or "wife"). For example, if a female appeared in the sample at each age between thirty-five and fifty-four, A_1 would be her average labor income at ages thirty-five through thirty-nine, and A_2 would be average labor income at ages fifty through fifty-four. A_1 is used to extrapolate her labor income backward—from age thirty-four through age twenty-three. Her imputed income at age thirty-four equals A_1 times the ratio of F_{34} to the average of F_{35} – F_{39} ; her imputed income at age thirty-three equals A_1 times the ratio of F_{33} to the average of F_{35} – F_{39} , and so on. A similar procedure is adopted when extrapolating individuals' labor incomes forward—beyond their oldest age as a PSID "head" or "wife" (fifty-four for the male in the current example) through age sixty-six.

The sample of individuals for whom we could compute/impute lifetime earnings profiles in this manner included 4,706 males and 5,278 females. However, our simulation includes only 2,000 individuals of each sex. To arrive at 2,000 male and female lifetime earnings profiles, we proceed in the following manner. First, we replicate each of the 4,706 male and 5,278 female observations in proportion to their sample population weights. Second, we group to the nearest integer the "blown-up" males and females into 0.05 percentiles of the present value of lifetime earnings. Third, we pick the observation with the highest value of lifetime earnings in each of the 0.05 percentiles. This gives us 2,000 males and 2,000 females whose level and pattern of lifetime earnings are used in the simulation.

3.3.9 Calibrating the Degree of Assortative Mating and the Inheritability of Skills

Armed with the original (not blown up by population weights) lifetime earnings of the sample of 4,706 males and 5,278 females, we formed for each year of the PSID sample the rank correlation coefficients between (a) the lifetime earnings of husbands and their wives observed in that year and (b) the lifetime earnings of fathers (mothers) and sons (daughters) observed in that year. The average (across all the years) values of the husband-wife and same-sex parent-child rank correlation coefficients are 0.10 and 0.44, respectively. These coefficients are used to calibrate the parameters mentioned above.

Before proceeding, we should point out that the 0.10 husband-wife rank correlation coefficient is quite small relative to our priors. It is also much smaller than the 0.70 guesstimate used by Gokhale et al. (2001). Although this estimate suggests that assortative mating on the basis of lifetime earnings may not be an important factor, one should bear in mind that the PSID underrepresents upper-income earners for whom assortative mating may be much more important.

3.3.10 Bequests and Inheritances

When a spouse dies, the surviving spouse retains all the couple's marital wealth. That is, all bequests of married agents go to their spouses. When a widow or widower dies, or when both spouses die at the same time, the wealth of the decedent(s) is evenly divided among the children.

3.3.11 Initial Wealth Endowments and Length of the Simulations

We start each of our simulations by giving all adults at $t = 0$ an endowment of wealth of one unit. We then run the model for enough years into the future until the distribution of wealth of sixty-seven-year-olds, as well as the total amount of wealth in the economy, stabilizes. Since the asymptotic wealth distribution as well as the total level of wealth is independent of the initial level and distribution of wealth, the fact that we start with this particular initial endowment of wealth does not alter our results. In practice, both the wealth distribution of sixty-seven-year-olds and the total level of wealth converge well before 150 years in each of our simulations. Nonetheless, to guarantee consistency across simulations, we run each simulation for 150 years.

3.3.12 Including Interest Rate Heterogeneity

Different households face different rates of return on their portfolios because they systematically choose to hold different portfolios. To incorporate rate-of-return heterogeneity, we use data on the portfolio holdings of households from the 1995 SCF. Our first step entails classifying those household assets reported in the survey into several categories. Next, we assign a rate of return to each asset category and compute the portfolio-weighted rate of return that each household faces, given its portfolio of assets.¹¹ Finally, we compute the weighted frequency of households for

11. The asset categories are liquid assets, government bonds, private bonds and bond mutual fund shares, stocks and stock mutual funds, real estate, and other nonfinancial assets. Liability categories include mortgage and real estate debt and other debt. In forming a weighted average rate of return on each household's portfolio, we used the absolute value of liabilities. For liquid assets, we assumed the geometric average annual real rate of return (0.68 percent) on U.S. Treasury bills during the period 1926–97. For government bonds, we use the geometric average annual real rate of return on long- and intermediate-term government bonds between 1926 and 1997 (2.09 percent). For private bonds and bond mutual funds, we use the geometric average rate of return on long-term corporate bonds between 1926 and 1997 (2.52 percent). For stocks and stock mutual funds, we use the weighted average of real rates of return on large and small company stocks (8.00 percent). The weights for the two stock market returns were obtained from analysts at the Wilshire 5000 company. The source for the aforementioned average rates of return is the 1998 Yearbook published by Ibbotson Associates. The average rates of interest on mortgages and other real estate debt were constructed using data from Case and Shiller (1990), who report annualized excess returns (excess over the three-month T-bill return) on home purchases for each quarter between 1971 and 1986 in four large U.S. metropolitan areas. We computed the total returns

rates of return ranging from 0 to 10 percent or more in steps of 0.5 percent. This frequency distribution is used to randomly allocate the average rate of return within each step to households in the simulation. Households are assumed to earn their assigned rate of return in each year of their lives.

3.3.13 Additional Issues of Calibration

The mortality probabilities used in the analysis are those released by the U.S. Social Security Administration for 1995. The interest rate used in the simulations is 4 percent. To obtain a realistic shape of the age-consumption profile, we assume that the time preference factors generate a 1.5 percent growth in the planned path of consumption per equivalent adult through age sixty-five.¹² From age sixty-six through eighty-eight, planned consumption per equivalent adult remains constant. As mentioned, our fertility matrix is derived from simulating CORSIM. Its fertility module includes separate logistic functions for thirty different subgroups of women estimated using data from the National Longitudinal Survey.

The subgroups are distinguished by age, the presence of children, marital status, race, and work status. The regressors in the logits are age, duration of current marriage, earnings, family income, homeowner status, marital status, schooling status, work status, and duration since the birth of the women's two youngest children. In producing the larger birth matrix from which we selected 2,000 rows, we ran the CORSIM model from its start year of 1960 through 2000. In so doing, we used the entire panoply of CORSIM modules to assign CORSIM agents the various socioeconomic characteristics, such as work status, entering as regressors in the fertility logits.

3.4 Findings

Tables 3.2 and 3.3 describe the distributions of wealth resulting from eleven different simulations. Our base-case simulation incorporates mortality prior to age eighty-eight, assortative mating, heterogeneity in skills, inheritance of skills, life-cycle consumption growth, heterogeneity in rates of return, progressive income taxation, and a social security system with

by adding the annualized real T-bill return for each quarter, then calculated the geometric mean over the period of the study and averaged the rates of return over the four metropolitan areas. This procedure yields a real rate of return of 0.45 percent. Finally, the rate of return for mortgage and real estate debt was calculated as the geometric average nominal mortgage rate between 1973 and 1997 divided by the geometric average rate of inflation over the same period. This yielded 3.91 percent. The average real rate on other debt was assumed to be 13.54 percent of the rate applicable for 1995, obtained from the Statistical Abstract for the United States, 1998, table 820.

12. Actual consumption per equivalent adult prior to age sixty-six will differ from planned consumption if the household receives one or more inheritances.

Table 3.2 Inequality and Bequest Flows: Base Case and Alternative Simulations

Case	Simulation	Wealth Gini	Consumption Gini	Bequest/Earnings	Cross-Generational Bequest/Earnings
1	Base case	0.336	0.283	4.5	1.5
2	No skill differences	0.097	0.045	4.3	1.4
3	No marital sorting	0.324	0.270	4.6	1.5
4	No inheritance of skills	0.327	0.277	4.6	1.5
5	No consumption growth	0.427	0.286	1.8	0.6
6	No interest rate heterogeneity	0.332	0.279	4.5	1.4
7	No progressive income taxes	0.436	0.317	3.2	1.0
8	No inheritances	0.325	0.287	0.0	0.0
9	No Social Security	0.278	0.271	10.6	3.4
10	No ceiling on SS taxable income	0.242	0.272	3.9	1.2
11	No inheritances and no Social Security	0.285	0.278	0.0	0.0

Source: Authors' calculations.

a ceiling on taxable income. The other ten simulations are variants on the base case that leave out one, or, in the case of simulation eleven, two of these elements.

The first row in table 3.2 shows results for our base-case simulation, including a wealth Gini of 0.336 and a consumption Gini of 0.283 among age-sixty-six households. The consumption Gini is smaller because consumption is financed in part by Social Security benefits, and these benefits are more evenly distributed than is net wealth. The first row also indicates that the steady-state flow of total bequests in our model is 4.5 percent of aggregate labor earnings, and the cross-generational flow is 1.5 percent of aggregate labor earnings. This cross-generational flow of bequests appears to be about half as large as in the actual U.S. economy. We say this because our own unpublished calculations using the 1962 and 1995 SCFs suggest that cross-generational bequests in the actual economy are roughly 3 percent of total labor compensation.¹³

Figure 3.1 graphs the first row of table 3.3—the wealth distribution of

13. In forming these calculations, we first benchmarked 1962 and 1995 SCF net worth and life insurance holdings to Federal Reserve Flow of Funds and American Council of Life Insurance respective reported aggregates. Next, we calculated the flow of bequests by age and sex by assuming that (a) single people with children leave their entire estates to their children, (b) married decedents leave 15 percent of their estates to their children, and (c) all decedents have the same age- and sex-specific mortality rates as those that prevailed in 1962 and 1994. Our estimated ratios of cross-generational bequests to total labor compensation are 0.030 for 1962 and 0.033 for 1994.

Table 3.3 Simulated Wealth Distributions

Case	Simulation	% of Wealth Held by Top									
		99	95	90	75	50	25	10	5	1	
1	Base case	100.0	99.2	97.4	89.8	71.8	47.6	27.6	18.4	7.0	
2	No skill differences	99.4	96.7	92.9	80.3	56.9	30.5	13.5	7.2	1.6	
3	No marital sorting	100.0	99.1	97.2	89.3	71.0	46.7	26.8	17.8	6.6	
4	No inheritance of skills	100.0	99.2	97.4	89.5	71.2	46.8	26.9	17.8	6.6	
5	No consumption growth	100.0	99.5	98.2	92.2	76.9	54.9	35.4	25.1	10.3	
6	No interest rate heterogeneity	100.0	99.3	97.5	89.7	71.6	47.1	27.1	17.9	6.5	
7	No progressive income taxes	100.0	99.5	98.4	92.8	78.1	55.5	34.4	23.8	9.4	
8	No inheritances	100.0	99.4	97.7	89.8	71.2	46.3	26.2	17.2	6.3	
9	No Social Security	100.0	99.2	97.4	89.1	69.0	42.1	21.2	12.6	3.8	
10	No ceiling on SS taxable income	100.0	99.1	97.0	87.8	66.8	39.4	18.9	10.7	2.9	
11	No inheritances and no Social Security	100.0	99.4	97.7	89.6	69.5	42.4	21.3	12.7	3.9	

Source: Authors' calculations.




Fig. 3.1 Comparison of wealth distributions: simulated, SCF 1995, and PSID 1984

age-sixty-six couples generated by our base-case simulation. It also compares this distribution with wealth distributions of married couples whose heads are aged sixty to sixty-nine in (a) the 1995 SCF and (b) the 1984 PSID. As the figure makes clear, the concentration of wealth at the very top of the wealth distribution is much more substantial in the SCF. The SCF wealth Gini is 0.73—more than twice as large as our simulated Gini.

Why is our model generating too little skewness in the distribution of wealth compared to the SCF? The answer, we believe, is that (a) the PSID distribution of lifetime earnings is much less skewed than that of the actual economy, and (b) the skewness of lifetime earnings makes a significant difference in the skewness in our simulated distribution of wealth. To see the importance of our assumed degree of earnings inequality to our simulated wealth inequality, consider the results in the second rows of tables 3.2 and 3.3, which omit labor earnings heterogeneity. Without skill differences, the wealth Gini falls from 0.336 to 0.097. The fact that wealth inequality exists in the absence of skill inequality is due to differences across households in the number of timing of their children as well as in the rates of return earned on their saving.

The PSID earnings data from which we construct our male and female distributions of lifetime earnings is certainly much less skewed than the corresponding male and female cross-section distributions of earnings in the SCF. This would explain why the simulated wealth distribution in our base case is less concentrated than in the SCF.

However, it is harder to explain why the PSID's wealth concentration in the top tail is much greater than that simulated by the model. This may




Fig. 3.2 Simulated and SCF-1995 distributions of wealth among bottom 50 percent of wealth holders

reflect a higher degree of underreporting of earnings among the rich than of underreporting of their assets. One way to test this possibility is to compare the distribution of wealth among the poorest (measured in terms of wealth) of our age-sixty-six simulated couples with the corresponding distribution of wealth among the poorest half of the SCF sixty- to sixty-nine-year-old couples. Figure 3.2 makes this comparison. While our model fits the distribution of the truncated data quite well at the very top of the tail, the fit in the rest of the top tail is not very good.

3.4.1 Determinants of Wealth Inequality

The remaining simulations in tables 3.2 and 3.3 examine other determinants of wealth inequality. Consider first the case 3 simulation, which eliminates assortative mating. Given the limited assortative mating we find in the data and incorporate in the base case, it is not surprising that eliminating assortative mating altogether makes very little difference to simulated wealth inequality. A second “nonfactor” with respect to the distribution of wealth is the inheritance of skills from parents. The source of agents’ skills does not appear to have much impact on the cross-section distribution of wealth. This is clear from the case 4 simulation, which eliminates the inheritability of skills, leaving the Gini at essentially its base-case value.

The case 5 simulation eliminates consumption growth, which means that it alters the model’s time-preference factors so that agents wish to have the same living standard per equivalent adult as they age. This smoothing of consumption generates more liquidity-constrained house-

holds that arrive at age sixty-six with zero wealth, which result, in turn, implies more wealth inequality. The wealth Gini in this case is 0.427—27 percent higher than in the base case.

The case 6 simulation turns off interest rate heterogeneity, which makes little difference to the distribution of wealth. The explanation here is that, other things being equal, households earning higher rates of return can afford to and do consume more at each point in time. Since accumulated assets reflect accumulated differences between past levels of earnings and consumption, this factor lowers the age-sixty-six wealth holdings of couples earning high rates of return. On the other hand, the smaller differences between past earnings and consumption are accumulated at a higher rate of return. This factor raises the age-sixty-six wealth holdings of high rate of return couples. According to the results, these two factors roughly cancel leaving wealth inequality essentially the same as in the case of no interest rate heterogeneity. In contrast to interest rate heterogeneity, progressive income taxation does play an important role in influencing wealth. In the case 7 simulation, this form of taxation is eliminated, raising the Gini coefficient from 0.336 to 0.436. Additionally, the share of wealth held by the top 1 percent of the wealth distribution rises from 7 percent to 9.4 percent.

3.4.2 The Role of Inheritances and Social Security in Wealth Inequality

How important are bequests and their associated inheritances to wealth inequality? The case 8 simulation addresses this question. It sets the probability of dying prior to age eighty-eight at zero. Compared to the base case, this experiment reduces wealth inequality only slightly. So, in the presence of Social Security, inheritances make the distribution of wealth more unequal. Is the same true if Social Security is absent? The answer, found by comparing cases 9 and 11, is no. With inheritances, but without Social Security, wealth inequality is somewhat smaller—the same finding reported in Gokhale et al. (2001).

How much does Social Security itself raise wealth inequality? Comparing the case 9 and 11 simulations reveals that the answer is a fair amount. Adding Social Security to the model's other features raises the wealth Gini by over 20 percent. It also almost doubles the share of wealth held by the top 1 percent of wealth holders. Figure 3.3 shows the impact of removing Social Security on the wealth levels of age-sixty-six couples at different percentiles of the wealth distributions. For those in the 75th percentile, eliminating Social Security would raise their age-sixty-six wealth by a factor of 2.55. In contrast, for those in the top 1 percentile, eliminating Social Security would raise their wealth by a factor of only 1.55.¹⁴

14. This difference would be greater still were we to assume that those who earn much higher than average earnings also earn much higher than average rates of return.




Fig. 3.3 Average wealth holding within selected percentiles of the age-66 wealth distribution, with and without Social Security

Why does Social Security increase wealth inequality? A small part of the reason, as previously mentioned, is that Social Security transforms inheritances into a disequalizing force. According to table 3.2, it also reduces the intergenerational flow of inheritances by over 50 percent. However, the main reason that Social Security is disequalizing is simply, as confirmed by the case 10 simulation, the ceiling that Social Security applies to its tax collection. This ceiling treats the lifetime rich more favorably than the lifetime poor. It also differentially annuitizes the lower classes. In the absence of the ceiling, the wealth Gini would be 0.242, compared to 0.336 in the base case, and the share of wealth held by the top 1 percent would be 2.9 percent, compared with 7 percent in the base case.

One can also assess the importance of Social Security to wealth inequality by considering how the wealth Gini varies with Social Security’s load factor. As mentioned, our base case assumes that, for every dollar workers pay in Social Security payroll taxes, they receive thirty cents in retirement benefits, when both the taxes and benefits are measured in present value. A modest reduction in this load, such that workers receive back, in present value, thirty-five cents on each dollar contributed, would raise the wealth Gini from 0.336 to 0.363—that is, by about 10 percent. It would also, incidentally, lower the flows of both total and cross-generational bequests by 13 percent.

3.5 Conclusion

Given our use of the PSID, which undersamples very high-earning Americans, to calibrate the distribution of lifetime earnings, it is not surprising that our model fails to reproduce the extreme skewness of the top tail of U.S. wealth distribution. On the other hand, our analysis confirms the proposition first advanced by Feldstein (1976) and reemphasized by Gokhale et al. (2001) that Social Security exacerbates wealth inequality by leaving the lifetime poor with proportionately less to save, less reason to save, and a larger share of their old-age resources in a nonbequeathable form than the lifetime rich. In so doing, Social Security denies the children of the poor the opportunity to receive inheritances. Consequently, inheritances become a cause of wealth inequality rather than wealth equality. All told, Social Security appears to be raising wealth inequality, as measured by the Gini coefficient, by roughly one-fifth, substantially increasing the share of total wealth held by the richest members of society, and greatly reducing the flow of bequests to the next generation.

References

- Abel, Andrew, and Laurence J. Kotlikoff. 1994. Intergenerational altruism and the effectiveness of fiscal policy: New tests based on cohort data. In *Savings and bequests*, ed. Toshiaki Tachibanaki, 167–96. Ann Arbor, Michigan: University of Michigan Press.
- Altonji, Joseph, Fumio Hayashi, and Laurence J. Kotlikoff. 1992. Is the extended family altruistically linked? Direct tests using micro data. *American Economic Review* 82 (5): 1177–98.
- . 1997. Parental altruism and inter vivos transfers: Theory and evidence. *Journal of Political Economy* 105 (6): 1121–66.
- Atkinson, Anthony B. 1971. The distribution of wealth and the individual life cycle. *Oxford Economic Papers* 23:239–54.
- . 1980. Inheritance and the redistribution of wealth. In *Public policy and the tax system*, ed. G. M. Heal and G. A. Hughes. London: Allen and Urwin.
- Atkinson, Anthony B., and A. J. Harrison. 1978. *Distribution of personal wealth in Britain*. Cambridge: Cambridge University Press.
- Atkinson, Anthony B., and Stiglitz, Joseph. 1980. *Lectures in public economics*. Singapore: McGraw-Hill.
- Auerbach, Alan J., Jagadeesh Gokhale, Laurence J. Kotlikoff, John Sabelhaus, and David Weil. 1995. The annuitization of Americans' resources: A cohort analysis. NBER Working Paper no. 5089. Cambridge, Mass.: National Bureau of Economic Research, April.
- Becker, Gary S., and Nigel Tomes. 1979. An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy* 87 (6): 1153–89.
- Blinder, Alan S. 1973. A model of inherited wealth. *Quarterly Journal of Economics* 87 (4): 608–26.

- Boskin, Michael J., and Laurence J. Kotlikoff. 1985. Public debt and U.S. saving: A new test of the neutrality hypothesis. *Carnegie-Rochester Conference Series on Public Policy* 23 (autumn): 55–86.
- Caldwell, Steven, Melissa Favreault, Alla Gantman, Jagadeesh Gokhale, Thomas Johnson, and Laurence J. Kotlikoff. 1999. Social Security's treatment of postwar Americans. In *Tax policy and the economy*, vol. 13, ed. James M. Poterba, 109–48. Cambridge, Mass.: MIT Press.
- Case, Karl E., and Robert J. Shiller. 1990. Forecasting prices and excess returns in the housing market. *AREUEA Journal* 18 (3): 253–73.
- Davies, James B. 1982. The relative impact of inheritance and other factors on economic inequality. *Quarterly Journal of Economics* 97 (3): 471–98.
- Feldstein, Martin. 1976. Social Security and the distribution of wealth. *Journal of the American Statistical Association* 71 (356): 800–07.
- Flemming, John S. 1976. On the assessment of the inequality of wealth. In *Selected evidence submitted to the Royal Commission: Report no. 1, initial report of the standing reference*, 34–70. London: HMSO and Royal Commission on the Distribution of Income and Wealth.
- . 1979. The effects of earnings inequality, imperfect capital markets, and dynastic altruism on the distribution of wealth in life cycle models. *Economica* 46:363–80.
- Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus. 1996. Understanding the postwar decline in United States saving: A cohort analysis. *Brookings Papers on Economic Activity* 1:315–90.
- Gokhale, Jagadeesh, Laurence J. Kotlikoff, James Sefton, and Martin Weale. 2001. Simulating the transmission of wealth inequality via bequests. *Journal of Public Economics* 79:93–128.
- Hall, Robert E. 1988. Intertemporal substitution in consumption. *Journal of Political Economy* 96 (2): 339–57.
- Harbury, C. D., and D. M. W. N. Hitchens. 1979. *Inheritance and wealth inequality in Britain*. London: George Allen & Unwin.
- Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff. 1996. Risk-sharing between and within families. *Econometrica* 64 (2): 261–94.
- Hurd, Michael D. 1990. Research on the elderly: Economic status, retirement and consumption & saving. *Journal of Economic Literature* 28:565–637.
- . 1992. Measuring the bequest motive: The effect of children on saving by the elderly in the United States. In *Savings and bequests*, ed. T. Tachibanaki, 111–36. Ann Arbor: University of Michigan Press.
- Laitner, John. 1979a. Household bequests, perfect expectations, and the national distribution of wealth. *Econometrica* 47 (5): 1175–93.
- . 1979b. Household bequest behaviour and the national distribution of wealth. *Review of Economic Studies* 46 (3): 467–83.
- Laitner, John, and F. Thomas Juster. 1996. New evidence on altruism: A study of TIAA-CREF retirees. *American Economic Review* 86 (4): 893–908.
- Loury, Glenn. 1981. Intergenerational transfers and the distribution of earnings. *Econometrica* 49 (4): 843–67.
- Meade, James E. 1964. *Efficiency, equality, and the ownership of property*. London: George Allen & Unwin.
- . 1966. Life-cycle savings, inheritance and economic growth. *Review of Economic Studies* 33:61–78.
- . 1976. *The just economy*. London: George Allen & Unwin.
- Menchick, P. 1979. Intergeneration transmission of inequality: An empirical study of wealth mobility. *Economica* 46:349–62.
- Miles, David. 1997. Demographics and saving: Can we reconcile the evidence? IFS Working Paper no. 97/6. London: Institute for Fiscal Studies.

- Oulton, N. 1976. Inheritance and the distribution of wealth. *Oxford Economic Papers* 28:86–101.
- Pryor, F. 1973. Simulation of the impact of social and economic institutions on the size distribution of income and wealth. *American Economic Review* 63:50–72.
- Stiglitz, Joseph E. 1969. Distribution of income and wealth among individuals. *Econometrica* 37 (3): 382–97.
- Tomes, Nigel. 1981. The family, inheritance, and the intergenerational transmission of inequality. *Journal of Political Economy* 89 (5): 928–58.
- Wilhelm, Mark O. 1996. Bequest behavior and the effect of heirs' earnings: Testing the altruistic model of bequests. *American Economic Review* 86 (4): 874–92.
- . 1997. Inheritance, steady-state consumption inequality, and the lifetime earnings process. *Manchester School of Economic and Social Studies* 65 (4): 466–76.
- Wolfson, M. 1977. *The causes of inequality in the distribution of wealth: A simulation analysis*. Ph.D. diss. Cambridge University, Cambridge, England.
- . 1979. The bequest process and causes of inequality in the distribution of wealth. In *Modelling the intergenerational transmission of wealth*, ed. J. D. Smith. New York: National Bureau of Economic Research.

Comment R. Glenn Hubbard

I enjoyed this chapter very much, and I believe the research program in which Jagadeesh Gokhale and Larry Kotlikoff are engaged will make significant contributions to our understanding of the link between social insurance design and the distribution of wealth. The chapter's purported focus is Social Security, but this is true only in a narrow sense. The present chapter is better seen as part of a broader research program by the authors to explain the distribution of wealth and determinants of wealth inequality (see especially Gokhale et al., forthcoming).

My comments are organized as follows. First, I pose some open questions to frame the analysis in the chapter. Second, I briefly review the chapter's principal findings. Third, I focus on promising extensions of the chapter in research and policy analysis.

Open Questions

In the background of the chapter lie two important questions. Why might Social Security affect the distribution of wealth, and what non-Social Security factors explain the actual distribution of wealth?

At least four factors offer a link between Social Security and the distribution of wealth. First, intergenerational transfers in a pay-as-you-go system may reduce aggregate saving and capital accumulation, and crowd

R. Glenn Hubbard is chairman of the president's Council of Economic Advisers, on leave from posts as the Russell L. Carson Professor of Economics and Finance at Columbia University and a research associate of the National Bureau of Economic Research.

out the private wealth accumulation of many households. Second, Social Security may affect private intergenerational transfers and bequests. Third, even a fully funded and actuarially fair Social Security system would reduce saving on account of imperfections in private annuity markets; households wishing to annuitize more of their household wealth may reduce private saving accordingly. Fourth, even a funded, actuarially fair system can affect the saving of many households in the presence of borrowing constraints (that is, when households cannot borrow against future benefits).

These channels raise the open questions of why households save and why some households save very little. Early analysis of Social Security and saving focused on effects of the program on different generations in a life-cycle setting (Feldstein 1974; Auerbach and Kotlikoff 1987). Within a generation, insurance-market and capital-market imperfections can connect Social Security (and social insurance more broadly) to the distribution of wealth (Abel 1985; Hubbard and Judd 1987; Hubbard, Skinner, and Zeldes 1995). Even relatively complicated versions of life-cycle models, however, have difficulty in explaining the saving of the wealthy.

Toward that end, life-cycle models' predictions often fail to match the highly skewed distribution of wealth in U.S. microeconomic data. Why the rich save so much over their lifetimes remains an open question, in spite of recent analyses of heterogeneity in rates of time preference, market imperfections and precautionary saving, fiscal incentives, and household variation in access to different rates of return.

As I argue below, understanding links between Social Security and household wealth accumulation and sources of heterogeneity in household saving rates are important for interpreting the chapter's results and policy implications.

Findings of the Chapter

Under fairly reasonable assumptions, the model used by Gokhale and Kotlikoff can replicate many (though not all) features of the real-world distribution of U.S. household wealth. As sources of wealth inequality, the chapter focuses on uncertain lifetimes, skill differences within and across generations, annuitization of Social Security "wealth," and skewness in the distribution of households' lifetime earnings. In particular, the authors argue that annuitization of Social Security wealth—an element of the current Social Security system, but not necessarily an element of proposals to introduce individual Social Security accounts—can worsen household wealth inequality.

The chapter uses the CORSIM microsimulation model of the U.S. economy (see Caldwell et al. 1999) with parameters based on calibrations using data from the Panel Study of Income Dynamics (PSID). The basic model has a life-cycle structure. Households have a stochastic lifetime of up to eighty-eight periods. There are an equal number of males and females (fer-

tility assumptions yield an equal match), with no population growth across cohorts. Marriage occurs at age twenty-two (the authors investigate both random and assortative mating).¹ In the base case, there are no differences in skill endowments; alternative cases consider contemporaneous skill differences (with and without assortative mating) and inherited skill differences.

Households' preferences are described by time-separable isoelastic functions of their own and future consumption, as well as that of their children through age twenty-two. While the basic structure used by the authors is standard among researchers employing life-cycle simulation models to study household consumption decisions, that structure begs deeper questions about saving decisions of very low- and very high-income households that may prove crucial for analysis of wealth inequality.

In their simulations, the authors assume a very low intertemporal elasticity of substitution in consumption, so that households effectively put great weight on their "safe" resources in deciding how much to consume at each point in time. The rate of interest is assumed to equal four percent per annum; time preference factors yield 1.5 percent annual growth in per capita consumption through age sixty-five. Mortality probabilities are taken from 1995 calculations by the Social Security Administration. Basic earnings and fertility matrices are taken from CORSIM. In general, the chapter might benefit from a bit more description of the calibration exercise and from sensitivity analysis; more discussion comparing the "earnings module" with other efforts using the PSID would also have been useful.²

The chapter's principal results are contained in tables 3.1 and 3.2, which report, respectively, for various simulations, Gini coefficients, and shares of total household wealth by wealth group. For the latter, the chapter also reports tabulations from the Federal Reserve's 1995 Survey of Consumer Finances.

Several features of the findings presented in tables 3.1 and 3.2 are partic-

1. Life doesn't have quite the feeling of Thornton Wilder's *Our Town* in the CORSIM model. Gokhale et al. (forthcoming), for example, describe marriage and death as follows: "22-year-old males and females were married to each other (at random, or according to their skill ranks, depending on the case being considered)," and "Oldsters were killed off at random according to the conditional probability of dying at their respective ages, and the existing wealth of those who just died was transferred to the surviving spouse (first) or children (equally)."

Previous simulation studies generally assumed that inheritances are unexpected and taxed away and redistributed (see, e.g., Hubbard and Judd, 1987; Hubbard, Skinner, and Zeldes, 1995; and Huggett 1996); this lack of emphasis on intentional bequests is consistent with the empirical work of Hurd (1992) and Gokhale, Kotlikoff, and Sabelhaus (1996). Previous work on intentional bequests and inheritability of skills (Flemming 1979) did not incorporate marriage.

2. Exploring differences between the present paper and the CORSIM approach in Gokhale et al. (forthcoming) might shed light on implications of key CORSIM assumptions. In the present paper based on the PSID, for example, assortative mating on the basis of lifetime earnings appears much less important than assumed by Gokhale et al. (forthcoming).

ularly noteworthy. While the predicted bequest flow in table 3.1 is significant, lifetime uncertainty per se (in the presence of Social Security) has a relatively modest effect on the wealth Gini; removing Social Security as well strengthens this effect.³ More significant effects are traceable to skill differences. Removing desired consumption growth implies that households want to maintain an equivalent living standard as they age; as a consequence of this consumption smoothing, more liquidity-constrained households approach retirement with no wealth, increasing wealth inequality. Interest rate heterogeneity does not appear to be an important determinant of the wealth Gini (I return below to whether heterogeneity in rates of return is correctly measured). While the results of table 3.2 come closer to matching the SCF wealth shares than predictions of life-cycle models with no intra-cohort wealth and income variation, the model still fails to capture wealth accumulation at the top.

The villain of the inequality play is Social Security. In part, this reflects Social Security's effect on the distribution of inheritances. That is, annuitization of Social Security benefits increases inequality because high-income households have substantial annuitized wealth. In the model, the differential effect of Social Security annuitization traces principally to the fact that the payroll tax ceiling benefits high-lifetime-income households (see case 10 in table 3.1).

Even within the structure of the model, these effects of Social Security on household variation in wealth accumulation may be overstated. The authors' assumption of a "pure tax" rate (i.e., tax rate net of that fund benefits to the household) of 67 percent substantially exceeds estimates produced, for example, by the Congressional Budget Office (Mariger 1999). More important, a portion of any pure tax represents the servicing of debt incurred to provide unfunded benefits to previous generation. Hence one would not associate the removal of this tax with Social Security tax (as long as promised benefits will be paid from Social Security payroll tax revenues).

Ultimately, the value of the authors' approach to analyzing the contribution of Social Security to the distribution of household wealth derives at least in part from the model's effectiveness in explaining the distribution of wealth. While removing Social Security reduces the share of wealth held by the top one percent of households from 7 percent to 3.8 percent, the actual share in the 1995 data cited by the authors exceeds 30 percent.⁴ This discrepancy is surely a topic for future research.

While the chapter's findings are thoughtful and suggestive, they seem to

3. In Gokhale et al. (forthcoming), the interaction of lifetime uncertainty and skill differences can be "leveling," in the sense that larger inheritances are predicted to be received by offspring with skill levels lower than those of their parents.

4. In Gokhale et al. (forthcoming), the authors attempt to address this problem by increasing the skewness of the distribution of lifetime earnings relative to that predicted by the PSID.

call for two extensions of the Gokhale-Kotlikoff research program—one on the “rich” and one on the “poor.” First, how can one model more precisely sources of skewness in earnings and rates of return? In the Survey of Consumer Finances and the PSID, “entrepreneurs” (owners of active businesses) figure prominently in the upper tail of the distribution of wealth (see, for example, Gentry and Hubbard 1997, 1999; Quadrini 1997). For such households, capital-market imperfections could link saving and investment for such households, so that entrepreneurial “selection” and “investment” decisions can explain high preretirement saving rates for these households and slow dissaving in retirement (Gentry and Hubbard 1999, 2000).⁵ Such considerations may be particularly important if wealth is required to exploit (business-related) skill differences, and they suggest the desirability of integrating determinants of household business saving in life-cycle simulation models.

Second, for low-lifetime-income households, how should one think about annuitization and welfare, as opposed to annuitization and saving? With imperfections in annuity markets, Social Security’s mandatory annuitization may make these households better off, while reducing their saving (Hubbard and Judd 1987). In addition, more research is needed to shed light on the saving decisions of low-lifetime-income households—in particular, the extent to which they would save for retirement in the absence of Social Security’s tax and mandatory annuitization elements.

While I believe these extensions are significant, I want to end where I began, complimenting Jagadeesh Gokhale and Larry Kotlikoff for their careful modeling and analysis of economic links between the structure of Social Security and the distribution of household wealth. Their chapter should be on the “recommended reading” list for serious students of Social Security reform.

References

- Abel, Andrew B. 1985. Precautionary saving and accidental bequests. *American Economic Review* 75 (December): 777–97.
- Auerbach, Alan J., and Laurence J. Kotlikoff. 1987. *Dynamic fiscal policy*. Cambridge: Cambridge University Press.
- Caldwell, Steven, Melissa Favreault, Alla Gantman, Jagadeesh Gokhale, Thomas Johnson, and Laurence J. Kotlikoff. 1999. Social Security’s treatment of postwar Americans. In *Tax policy and the economy*, vol. 13, ed. James M. Poterba. Cambridge: MIT Press.
- Feldstein, Martin. 1974. Social Security, induced retirement, and aggregate capital accumulation. *Journal of Political Economy* 82 (September–October): 75–95.
- Flemming, John S. 1979. The effects of earnings inequality, imperfect capital mar-

5. Recent research (Quadrini 1999; Quadrini and Rios-Rull 1997) indicates that such descriptions of entrepreneurial saving and investment are critical for explaining the share of wealth held by wealthy households in general equilibrium.

- kets, and dynastic altruism on the distribution of wealth in life cycle models. *Economica* 46:363–84.
- Gentry, William M., and R. Glenn Hubbard. 1997. Distributional implications of introducing a broad-based consumption tax. In *Tax policy and the economy*, vol. 11, ed. James M. Poterba, Cambridge: MIT Press.
- . 1999. Entrepreneurship and household saving. Columbia University, Graduate School of Business. Mimeo graph.
- . 2000. Tax policy and entrepreneurial entry. *American Economic Review* 90 (May):
- Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus. 1996. Understanding the postwar decline in United States saving: A cohort analysis. *Brookings Papers on Economic Activity*.
- Gokhale, Jagadeesh, Laurence J. Kotlikoff, James Sefton, and Martin Weale. Forthcoming. Simulating the transmission of wealth inequality via bequests. *Journal of Public Economics*.
- Hubbard, R. Glenn, and Kenneth Judd. 1987. Social security and individual welfare: Precautionary saving, borrowing constraints, and the payroll tax. *American Economic Review* 77 (September): 630–46.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. Precautionary saving and social insurance. *Journal of Political Economy* 105 (April): 360–99.
- Huggett, Mark. 1996. Wealth distribution in life cycle economies. *Journal of Monetary Economics* 38:469–94.
- Hurd, Michael D. 1992. Measuring the bequest motive: The effect of children on saving by the elderly in the United States. In *Savings and bequests*, ed. T. Tuchi-banaki, Ann Arbor: University of Michigan Press.
- Mariger, Randall. 1999. Social security reform: A selective survey of the issues. Washington, D.C.: Congressional Budget Office, August.
- Quadrini, Vincenzo. 1999. Entrepreneurship, saving, and social mobility. *Review of Income and Wealth* 45 (March): 1–19.
- Quadrini, Vincenzo, and Victor Rios-Rull. 1997. Understanding the U.S. distribution of wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 21 (Spring): 22–36.

Discussion Summary

Jeffrey B. Liebman was interested in the effects of an investment-based system on bequests. In such a system with preretirement bequests, individuals from poor families will inherit more as a consequence of income related mortality. In addition, to the extent that individual accounts cause people with higher incomes to reduce other asset holdings resulting in little net new financial wealth, while lower-income individuals lack financial assets to substitute out of, individual accounts might make the wealth distribution more equal.

Martin Feldstein mentioned past work in which it was difficult to reconcile the existing financial wealth distribution with standard life-cycle saving models and the income distribution. However, if Social Security wealth was included with ordinary financial wealth, then the total wealth distribu-

tion was much more consistent with life-cycle saving models. If wealth is to be a measure of future consumption potential, then the present value of Social Security wealth should be included. However, liquid wealth provides options that an annuity (which may not have begun making payments) does not provide. For people who have not retired, wealth provides the option to take time when searching for a new employer and the ability to start a business or help one's children start a business. The present value of an annuity does not offer any of these opportunities. Therefore, it may be legitimate to focus on liquid wealth, especially when looking at younger people.

Stephen C. Goss did not find it particularly surprising that abolishing the current Social Security system has a dramatic effect on wealth when it is assumed that Social Security contributions do not create any wealth. He questioned whether it is appropriate to ignore Social Security wealth completely. In response, the authors said that when people express concern about wealth inequality, they are usually not talking about Social Security wealth. If social policy is to address this important issue, then it is important to examine wealth inequality in the absence of Social Security wealth.

Angus Deaton felt that while the PSID is not perfect, the SCF data set should not be considered the gold standard. Because the wealth distribution is very skewed, over-sampling the top end may help estimate the mean, but the upper percentiles will still be subject to a large amount of estimation error. For most models the wealth levels at the top of the distribution are not important and this is the data missing from the PSID. Furthermore, it is not necessary to build exotic behavior into simulations to generate large amounts of wealth inequality. The standard life-cycle model with earnings uncertainty implies that wealth inequality at the date of retirement is a very large number. Kotlikoff expressed doubt that a standard life-cycle model with earnings uncertainty could generate a Gini coefficient of 0.7.

R. Glenn Hubbard said that if the goal of the paper is to compare the poor to the middle class, then the approach taken in these simulations is fine; however, if the goal is to examine inequality for the entire wealth distribution, then it is important for the simulations to accurately reflect the behavior of the very wealthy—a group for whom the data are not very accurate.

Eytan Sheshinski emphasized the link between large unintended bequests and a lack of annuitization. He said it is a puzzle that middle and upper class people convert such a small percentage of their wealth into annuities.