

1

Demographics, Market Failure, and Social Security

Jerry R. Green

This chapter is an analysis of the argument for Social Security on the grounds that it can provide a form of insurance that is otherwise unavailable. Individuals in our society face many risks. Some of these are insurable; others, though not completely insurable in a contractual sense, can be mitigated by the cooperative behavior of families and benevolent organizations; still others can be offset by the guarantees implicitly or explicitly provided by government policy.

It is the last category that concerns us here. Within this is a certain class of risk that cannot be diversified over the members of a single generation. To mitigate against this type of risk, a mutual insurance arrangement or an implicit contract involving members of different generations is required. One function of the Social Security system might be to provide a framework in which this type of implicit contract can operate. The contracting parties are, in effect, the current generations and the as yet unborn.

Risks requiring this type of intergenerational insurance are of two types. First are the real economic risks associated with productivity fluctuations, technical progress, wars, and, in an international context, trading relations with the rest of the world. Second is a risk associated with demographic fluctuations, which are a risk to present generations because they affect the future of capital markets and hence the rates of return that can be earned on savings invested for retirement income. These fluctuations are a risk to future generations because, by their very nature, they affect the conditions of the labor markets in which people will operate.

Risks due to fluctuations in population growth require careful cost-benefit analyses. The set of people affected by economic policies is itself a variable. Policies that depend on the realized rates of growth will affect different sets of individuals in different circumstances. Therefore, the benefits of the policy cannot be properly measured by summing up the net willingness to pay over the population. Some of the population will be present only in some circumstances. Hence, we must develop a method for assessing the risk-reducing benefits of policies that affect individuals who are members of the population less than 100 percent of the time.

In the next section, I discuss the principal economic risks being considered. Pay-as-you-go policies, such as the social security system, will be modeled as intergenerational transfers, possibly contingent on the rate of population growth. The sum of all transfers must be zero at all points in time. Direct capital accumulation or decumulation by the government will not be considered. The transfers can be executed in a lump-sum fashion, without distorting market prices.

Following this, I present a set of ballpark estimates of the benefits from risk reduction that might be expected from such intergenerational transfer policies that are made contingent on the realized rate of population growth. I then consider the role of social insurance in this regard, using a two-period model of the individuals' lifetimes. By employing a set of models based on a more than two-period lifetime, I extend this model to include a more complex view of the individuals' saving decisions and a more detailed view of the possibilities for intergenerational transfers. I then give the mathematical details of the above results and provide a summary of these ideas.

Risks: Insurable and Uninsurable

In this section, I enumerate some of the risks that might properly be considered the targets of social insurance. Many economic risks that individuals and families face are not the proper subject of pay-as-you-go social insurance schemes. That is not to say that they are not the proper target for public policy more generally. Nevertheless, I omit them from further consideration.

We can categorize the risks that can be addressed by social insurance into three groups. First, we have risks that are common to all members of a given cohort; second, we have those that are common to all individuals living at the same time but are not necessarily confined to individuals of roughly the same age; finally, we have those that are of a permanent nature, affecting all individuals from some time forward.

The most important source of risk to each cohort is its own size. The larger a cohort is, the lower will be the wages that a typical member of the cohort will receive. As the cohort ages and begins to garner investment income on its accumulated savings, the level of the aggregate stock will increase and the gross return to additional capital accumulation will decrease. Savings behavior is itself determined by the interaction of individuals' motives for saving and their expectations about their needs and about current and future returns. The variability of the size of each generation is the principal source of risk studied in this chapter. This risk is of the first type mentioned above. All individuals of a similar age will be similarly affected. They cannot, therefore, insure themselves. They must pool their common risk with others whose risks are different, and that means pooling with those of different age groups.

The second category, risks common to all individuals contemporaneously alive, include short-run risks of productivity fluctuations, which might affect both the real wage and the return to capital, and other exogenous sources of risk. The need for the government sector to use resources for collective purposes such as defense are in this category. Similarly, risks of an international nature, such as exchange rates or other countries' commercial policies, are a source of economic fluctuation common to all citizens of our country, regardless of age. That is not to say that all individuals are affected by these risks in exactly the same way. There will always be a varying impact, and we should separate the total risk into intracohort and intercohort components.

The final category of risks includes those that tend to be permanent changes in the situation. They affect all current cohorts and are expected to continue to affect all individuals in the future. Examples of particular importance to social insurance are the lengthening of expected lifetimes; the changing health status of older citizens, which affects their need for income in their retirement years; and longer term social trends, such as the reduced frequency of elderly parents living with their children. In a certain sense, these can be viewed as changes in taste, but they are the type of changes that society is likely to validate by using the social insurance system to help spread the burden of the transition over more than one generation.

Another category of risks includes those of a personal nature: whether or not one will marry, have children, become a widow or widower, and so on. As important as these are at the individual level, they are not the proper subject of social insurance policy as considered in this chapter. It is essential, however, to remember that any social security system might wish to take an individual's personal status into account in determining benefits. And as Boskin and Puffert (1987) have shown, this can introduce substantial risks into the system because of the variability of benefit levels.

Rough Estimates of the Benefit of Social Insurance Policies to Mitigate Risks of Demographic Fluctuations

In this section, I give, in very approximate terms, the likely benefits of a policy that can reduce the risks of population fluctuations both for current workers, who will be retired later, and future workers. To begin this calculation, we must have some feel for the range in which the working population could vary. For concreteness, let's focus on a particular date in the future, say 2025.

Fertility rates in the United States have ranged from a high of 3.77 births (estimated total number of births per woman, throughout her childbearing years) to a low of 1.75 births. The high end of this range is surely a phenomenon related to World War II and the large immigration waves of the early

twentieth century. For these reasons, it is quite unlikely that this fertility level will be repeated again, and surely not before the year 2025.

A more plausible range is between 2.7 and 1.7 births. In table 1-1, I present the projections for the U.S. population under three alternative assumptions: 2.1 births which is approximately the rate necessary to maintain a constant population (in the absence of immigration), 2.7 births, and 1.7 births per woman.

Next we have to compute how the capital-labor ratio will vary as a result of these variations in population size. If the amount of saving is invariant to the realization of population growth rates, the capital-labor ratio will vary over a range that is 33 percent of its average value. Using a Cobb-Douglas production function with capital's share set at 25 percent, the ratio of the marginal products of capital would vary by about 33 percent.

This seemingly modest variation in the rental price of capital would produce very large swings in the total return to savings, accumulated over a holding period of say twenty years. Depending on our assumptions about the rate at which capital income is taxed, the real net rate of return (annualized) could be as low as 1 percent per year if population growth is slow, or as high as 2 percent per year if population growth is rapid. This translates into a 22 percent fluctuation in the real wealth accumulated for retirement due solely to the variation in the interest rates induced by alternative assumptions about the rate of population growth.

This range is further magnified when we look at the implications for retirement consumption. When accumulation has been rapid due to high real interest rates, retirement consumption will be high for the same reason. Although changes in fertility trends might partially offset these effects during the retirement years themselves, there will definitely be a positive correlation between accumulated savings and the retirement consumption they can finance. In some simulations I have done, which are not reported here, I have found that the average level of consumption feasibly could vary over the retirement interval 2025 to 2045 by as much as 34 percent, with similar fluctuations for other periods.

It is reasonable to believe that future retirees would be willing to take a 5 percent lower average consumption level to offset this risk to a significant extent, if possible, by means of social insurance. Thus, 5 percent of the retirement income of the generation that will be retired in the 2025 to 2045 interval (today's thirty-year-olds) is a very rough approximation of the benefit they would receive from the risk-reducing aspects of a social insurance policy that is contingent on the future rates of population growth.

To translate this into a present value, again very roughly, we can take the retirement income of this twenty-year cohort out of their accumulated savings to be in the range of \$1 trillion annually. Thus, this amount of risk reduction will be worth around \$50 billion (1986) in each year of the 2025

Table 1-1
Projection of the U.S. Population by Broad Age Groups

Year	Population (in thousands) as of July 1				65 and Over		Ratio of Persons Under 18 to Persons 18-64
	Under 18	18-64	65 and Over	Total	Percent of Total	Ratio to 18-64	
Assumption A—2.1 fertility							
1975	66,273	124,847	22,330	213,450	10.5	0.18	0.53
2000	71,079	160,815	30,600	262,494	11.7	0.19	0.44
2025	74,857	176,751	48,105	299,713	16.1	0.27	0.43
2050	78,701	188,448	51,247	318,396	16.1	0.27	0.42
Assumption B—2.7 fertility							
1975	66,273	124,847	22,330	213,450	10.5	0.18	0.53
2000	91,152	165,255	30,600	287,007	10.7	0.19	0.55
2025	121,054	212,852	48,105	382,011	12.6	0.23	0.57
2050	158,987	283,767	56,575	499,329	11.3	0.20	0.56
Assumption C—1.7 fertility							
1975	66,273	124,847	22,330	213,450	10.5	0.18	0.53
2000	57,322	157,176	30,600	245,098	12.5	0.19	0.36
2025	49,938	152,378	48,105	250,421	19.2	0.32	0.33
2050	44,146	135,532	47,024	226,702	20.7	0.35	0.33

Source: U.S. Bureau of Census, Current Population Report, Series P-25, No. 601, "Projections of Population of the United States: 1975-2050."

to 2045 retirement period. Discounting this to the present at a social discount rate of 2 percent per year, we have a present value of the policy, for this cohort alone, of \$19 billion annually.

A similar calculation for the generation that will be working in these same years reveals an equally substantial gain to be made from risk-reducing social insurance policies. The 22 percent fluctuation in the capital stock, when combined with the correlated influence of their own variable numbers, gives a range for the capital-labor ratio of about 29 percent. This produces a 22 percent variation in the real wage, using the same Cobb-Douglas assumptions as above. Again, we can estimate the amount of consumption that these workers, most of whom are not yet born, would be willing to forgo to reduce this risk substantially. A risk premium of 2 percent of their labor income is not unreasonable. There will be approximately 2 million households in each year's cohort, using the average assumptions about population growth. If each of these households has a real income of \$40,000 (1986), this means that they would be willing to pay, in aggregate, \$1.6 billion to avoid the risk. Thus, with an average working life of forty-six years, the entire working population would be willing to reduce their income by almost \$60 billion annually to substantially reduce the risk we are considering. Again, discounting this to the present, the present value of this policy of population-contingent income transfers is about \$23 billion annually.

Needless to say, the calculations above are very crude. The aggregate possible benefit in present value terms, restricting our attention to the 2025 to 2045 interval, is \$42 billion per year. This is a substantial sum. The purpose of this chapter is to see whether a population-contingent intergenerational transfer policy could be designed to capture at least part of these potentially large gains, and whether the gains that are captured can be spread equitably across individuals so that all would agree that the policy has been beneficial to them.

I cannot imagine that any policy could reduce the risks of population fluctuations to zero. Nevertheless, it is the nature of risk analysis that the most benefit is to be gained from mitigating the most extreme aspects of the variation. Based on the calculations above, it would seem that there are gains to be had in the multibillion-dollar-per-year range, if even part of the risk could be insured. Thus, we now turn to the more precise economic models in which the feasibility of a risk-reduction policy can be evaluated.

Mitigating the Risks of Demographic Fluctuations through Social Insurance

Two-Period Lifetime

In this section, I consider the principal risk for which social insurance is particularly well suited—the risk of fluctuations in the population size as caused

by changes in the birthrate. Separate considerations are necessary for individuals who are already alive and for those in cohorts to come because, as we shall see, the risks they face are different.

An individual in a previous cohort has certain expectations about the prices and interest rates that will prevail in the future, conditional on the population size. If the next cohort, whose size is at present uncertain, turns out to be larger than expected, the older cohort will experience an increase in the real return on its savings as soon as the younger cohort joins the work force.

Now let's imagine ourselves as individuals in the younger cohort. If our generation is larger than expected, we will earn lower wages than expected throughout our lifetime. The return on a unit of retirement savings might be higher or lower, depending on how elastic the supply of savings is with respect to the wage during our working years and with respect to the interest rate. Whether interest rates in our retirement years rise or fall, we will surely be worse off by having been born into a larger cohort.

Thus, individuals in younger and older cohorts have negatively correlated risks with respect to variation in the size of the cohort in question. Negatively correlated risks create a *prima facie* case for mutually beneficial insurance. Moreover, under my assumption that individuals in the same cohort must act simultaneously, the private sector cannot pool these risks effectively, for the younger agent, who should be a party to this insurance, cannot contract for it until he or she and all the members of his or her cohort are born. And once the size of the cohort is known, it is too late to insure this risk, as the relevant uncertainty has been eliminated.

In this chapter, I examine whether some social insurance contract might be mutually beneficial to present and future cohorts. A contract of this form would bind the as yet unborn to a system of income transfers contingent on the size of their cohort. In any realization, the members of the younger generation would be either taxed or subsidized by these transfers when they are young. Thus, a large generation would be subsidized when young because the invested capital of their elders would have been highly profitable and the elders would be subject to taxation on this basis. Later in life, they would be on the other side of the same contract. There would be a new risk to insure: the size of the next generation to come.

Let's first look at the situation in the absence of a social insurance system and under the further simplifying assumption of a two-period lifetime, with working years in the first period and retirement in the second. Population growth is random. In an equilibrium of this model, a typical generation's savings depend on the amount of capital in place when it is born and its size relative to the size of the immediately preceding cohort. The evolution of the capital stock over different generations can be obtained by applying this savings policy function to the distribution of successive generation sizes.

On this *laissez-faire* system we superimpose a social insurance scheme as follows. A pay-as-you-go system is one in which each retiree receives a

transfer from the government that is financed by a tax on the workers in an amount dependent on the size of their respective generations, and perhaps on the size of the capital stock. This transfer may be either positive or negative.

Given the considerations mentioned above, it is natural to consider transfers to retirees that depend negatively on the size of the succeeding generation. Let's suppose that the typical member of the retired generation is willing to accept the transfer, prospectively, before learning the size of the next generation. Would the unborn generation be willing to accept the other side of this gamble were it able to express its desires? If so, our argument is that a social insurance system that can commit future generations to such a policy is mutually beneficial to all concerned, even though in actuality some generations will be subsidized and others taxed.

The answer depends on the way in which we treat the members of the future generations. Imagine an individual who knows that he or she will surely be born as a member of the next generation but is unsure of exactly how many others will be born. His or her expected utility is determined by averaging over all possible population sizes.

In contrast to this, imagine an individual who does not know whether he or she will be born. His or her expected utility should be calculated conditional on this event. Thus, for an individual who will be born only if the rate of population growth is very high, expected utility will depend on the circumstances in which the social insurance scheme as envisioned would pay him a subsidy. Such individuals surely would prefer that this type of social insurance scheme be implemented. Conversely, under this view of the world, the individuals least likely to benefit from social insurance are those most likely to be born. To be universally beneficial, a social insurance scheme would have to keep these individuals at least as well off as they would be in the *laissez-faire* state.

If a social insurance system could be found that would pass this test and keep the older generation at least as well off, in expectation, it would benefit all. The individuals who are less likely to be born would receive more positive benefits. By being "absent" sometimes, they are "missing" in situations in which they would have been taxed. Thus, conditional on being born, they are doing better than they would be under *laissez-faire*.

Another approach is to treat all individuals as equally likely to be born. The actual composition of the population would then be determined in two stages. First, the actual size of the population would be fixed, but we would still be unsure about who, among all the potential individuals, will be present in the population. Second, the set of individuals comprising the population would be determined. Modeling random population growth in this way allows us to have a variable population size and individuals who are *ex ante* identical, without running afoul of the law of large numbers.

Under this set of hypotheses, only a single calculation need be made,

because all individuals in the same cohort are identical at the time our policy evaluation is made. Each individual assesses the possible worlds into which he or she might be born by forming a posterior belief about the size of his or her generation, given that he or she will be a member of it. The average population size computed in this fashion is larger than the *ex ante* mathematical expectation of the distribution of population sizes because of the positive correlation of the births across potential people. Since the large cohorts are the ones that are to be subsidized, the typical individual will assign more weight to them in his or her probabilistic judgment if he or she forms conditional beliefs in this way, as compared to the estimate of an outside observer who is not conditioning his or her expectations on the event of his or her own birth. Therefore, under this approach the use of social insurance to mitigate the risk of population fluctuations is more likely to succeed than it would be under the hypothesis that individuals' births are ordered.

In the ordered births model, one can improve the welfare of the retirees and an individual who will surely be born into the succeeding generation if the risks they face are sufficiently severe relative to the surplus that will, of necessity, go to the less likely to be born members of the latter's generation. The conditions for this are developed mathematically later in this chapter. They turn on the relationship between the individual's risk aversion and the elasticity of substitution between labor and capital in production. The less substitutable are the factors of production, the more their relative prices will have to vary in the equilibrium in order to absorb the inelastic supply of labor that will be present at that time. Thus, social insurance will be universally beneficial if the coefficient of risk aversion and the elasticity of substitution are both high. Quite surprisingly, perhaps, these conditions for mutually beneficial social insurance will not hold under the commonly estimated values for these parameters. Therefore, although this type of social insurance might be thought to be socially beneficial when its effects on all individuals in the population are taken into account, it is unlikely to be beneficial to everyone.

This rather negative conclusion is partially a by-product of the assumption of a two-period lifetime. In the next section, I will show that under more general specifications of the demographic structure, there is a greater chance for universally beneficial social insurance to mitigate the risks of population fluctuations, even under the ordered births hypothesis.

Under the symmetric births hypothesis posited by de Bartholome and Brandts (1986), there is a much better outcome, even in the two-period lifetime model. Social insurance provides the individuals in question with an actuarially fair bet that is perfectly negatively correlated with the risk they face. The reason that this bet is exactly actuarially fair, whereas it is at less-than-fair odds in the ordered births specification, is that the members of the younger generation put more probability weight on the event that they will be born into a larger cohort than do the typical members of the older cohort. This is

not because of a difference in beliefs or information. It is due simply to the fact that the younger agents neglect all those circumstances in which they themselves are not born, while the older agents do not care *who* is born, just *how many* are born.

More Than Two-Period Lifetime

In this section, I reconsider the above results in cases where agents live for more than two periods. If interpreted in economic terms, a two-period model ignores the part of life before one works and divides the remainder into two equal intervals. This distorts the analysis in several respects. A more realistic model must allow for the difference in length between the working and retirement periods and account for the first twenty or so years of life, which fall into neither category. Although theoretically interesting, it is not of practical relevance unless we can determine that the analysis is robust when the number of periods is extended.

Three-Period Model. Let's begin with a three-period model in which the working part of an individual's lifetime is modeled as being twice as long as retirement—work in the first two periods and retirement in the last. Thus, the ages involved are roughly 20–39, 40–59, and 60–80. Childhood and education are still neglected.

The economic difference between this and the two-period model is twofold. First, it allows for two periods within which each generation coexists with its immediate successor and its immediate predecessor, rather than the one period in the model discussed above. Second, because of this, adjacent generations can base the transfers between them on the size of the generation that follows them both. When this generation's size is known, the two contracting generations are at different times of life—one is retired, and the other is still working. Thus, their interests are opposed with respect to changes in the subsequent cohort's size. The retirees are hoping that they are more numerous so that the capital-labor ratio is lower, and the workers are hoping that they are less numerous so as not to depress their own wages. This effect is partially offset by the fact that the retirement income of the original workers will be higher if the generation following them is larger. Overall, however, it can be shown that this effect is overridden by the effect on the initial period's wages; the interests of the two generations in question are indeed opposed.

Under these conditions, there will be a mutually advantageous social insurance scheme if the retirees receive transfers from the immediately following generation. This will remain true under either of the assumptions about how the size of subsequent generations is determined, for unlike the case of the two-period lifetime, there is no uncertainty about the size of one's own generation. All the uncertainty pertains to a risk that is external to the two generations that are parties in the contract.

A contract between the retirees and the immediately following generation of workers would require age-specific taxation of the working population. The youngest generation of workers is, in this scheme, left out. If all workers have to share the burden of taxation because the older workers cannot be taxed or subsidized without the same being done to the youngest group, then this youngest group is placed in a position similar to the younger generation in the two-period lifetime case and my conclusions in the previous section still apply.

In summary, if the transfer scheme can distinguish between individuals of different generations, then a mutually beneficial social insurance system can be designed for all generations regardless of the modeling of demographic uncertainty. However, if the transfer policy can distinguish only between working-age and retirement-age individuals, it can benefit all generations only under the assumption that all individuals are *ex ante* identical.

Four-Period Model. The idea behind a four-period lifetime model is not simply that four is one more than three and that it is good to generalize the model further. The reason for exploring this model is that the three-year model omits the first twenty years of life. Even though individuals are not economically active in these years, they are important because accounting for them eliminates uncertainty about the size of the next working generation, even though there might still be uncertainty about the rate of population growth. The four periods are a period of economic inactivity, which serves only to model the size of future generations, then two working periods, and finally a retirement period of half the length of the working interval.

The results of the four-period model are qualitatively similar to those of the three-period model. Let's label the generations according to the period in which they are born and imagine that we are in period 4, just before the size of generation 4 is announced. Since these individuals will not be working, this announcement will have no impact on the capital-labor ratio or on the rate of return or the wage in the present period. Therefore, generation 1, which is now retired, experiences no real uncertainty on this account. Any transfer policy involving generation 1 that is conditional on the size of generation 4 will be welfare reducing for generation 1, unless it is more than actuarially fair. Therefore, the hope for a policy that is capable of being [Pareto improving] will have to rely on generations 2 and 3.

Generations 2 and 3 are both working at present. Their interests are opposed, however, because generation 3 will be retired when generation 1 starts working, whereas generation 2 will not. Thus, in a four-period model, mutually beneficial insurance is possible even though the size of one future generation is known in advance of its entering the work force because the working part of one's lifetime is twice the length of the retirement or the pre-working period. Note that this would not be true in a three-period model with all three phases of life of equal length. The retirees and the workers face no

uncertainty about the capital-labor ratio during their lifetime. The uncertainty that any one generation faces is present only during its first period of life. But at that stage members of the cohort are not yet economically active, and they have no income of their own. We assume that they are sharing in their parents' consumption without effectively diminishing it.

Mathematical Analysis

Two-Period Lifetime: Ordered Births Hypothesis¹

Let the initial capital stock be K_0 and the consumption of the current young generation (generation 0) be c_{y0} . Their savings for their retirement are equal to their wages, w_0 , minus this current consumption. If n_1 is the population growth factor, then their retirement consumption is

$$c_{r0}(n_1) = (w_0 - c_{y0})(1 + f')\left(\frac{w_0 - c_{y0}}{n_1}\right) \quad (1.1)$$

where f is the production function, as it depends on the capital-labor ratio. We assume that c_{y0} is chosen to maximize the agent's lifetime expected utility, holding other people's saving behavior constant and regarding n_1 as a random variable.

A population-contingent, pay-as-you-go, intergenerational transfer scheme is modeled as a term, $b(n_1)$, to be added to equation 1.1. If it is beneficial to all members of generation 0, we must have:

$$\sum_{n_1} u'_r(c_{r0}(n_1)) b(n_1) > 0 \quad (1.2)$$

where u_r is their retirement period utility function.

Now we look at the members of generation 1. Their lifetime utility is affected by the social insurance scheme in two ways. First, they must pay for generation 0's transfer (recall that this may be positive or negative). Second, the result of this transfer affects their own savings and hence their own retirement consumption, even in the absence of transfers between themselves and generation 2. The magnitude of this second effect depends on the sensitivity of future interest rates to the capital-labor ratio. As savings are predetermined, the capital-labor ratio depends only on n_1 . The magnitude of this dependence is a function of the elasticity of substitution of the production function, f .

The higher this elasticity, the less dependent is the future interest rate and therefore the less risky is retirement consumption as it is influenced by the

first period's transfer. Thus, generation 1 will benefit from the transfer if they are very risk tolerant and if the elasticity of substitution is high.

Specifically, if

$$1 < \frac{f'(f - n_1 f')}{f\sigma(1 + f')} R_r + \frac{k_1 f'}{f\sigma} R_y \quad (1.3)$$

where R_r and R_y are coefficients of relative risk aversion in the retirement and young periods, then the worst treated member of generation 1 will benefit from a social insurance scheme for which equation 1.2 holds with equality (i.e., generation 0 will accept it as well).

Examining equation 1.3 in more detail, we see that it is quite unlikely to hold. When utility is logarithmic, $R_y = R_r = 1$. And if production is Cobb-Douglas, the entire right side of equation 1.3 reduces to capital's share in this production function, which is necessarily less than unity. Hence, equation 1.3 is surely violated in this case.

More Than Two-Period Lifetime

In the more than two-period lifetime models developed earlier in this chapter, the parties are generations whose size is known when the random event upon which they are contracting is determined. Therefore, we know from the basic theory of behavior in the presence of risk² that there exists a sufficiently small gamble, which, if negatively related to exogenous random wealths, will be acceptable to both sides. The future savings behavior of the labor generation might be affected by these transfers. To the extent that savings increase and the elasticity of substitution is small, future interest rates, and hence their own retirement consumption, will be lowered. However, under typical conditions, this interest rate effect will be dominated by the direct risk-reduction benefits of the population-contingent transfer.

Note, however, that to implement this plan, generations 0 and 1 must be able to make a contract that excludes the members of generation 2, even though in the three-period model generation 2 is working, just as generation 1 is, in its first period of life. If we were to accept as a constraint on the social insurance compact that taxes or subsidies be common across workers and retirees, not only constant across members of the same generation, then the same set of considerations that made the two-period model unlikely to succeed for members of the youngest cohort would apply here as well.

Summary

This chapter has addressed three questions: Is the uncertainty inherent in the rate of population growth a serious source of inefficiency that is worth overcoming, if possible, by a system of contingent intergenerational transfers? Can these benefits actually be realized by a pay-as-you-go system? And how are the gains distributed over the population and across successive generations?

I have shown that the risks are substantial, though silent. My best estimate of the annual value of eliminating the income risks attendant to population fluctuations is \$42 billion. Of this, \$19 billion is the value to future retirees, and \$23 billion is the value to future workers.

I also have shown that many of these benefits can be captured by a system of implicit, or governmentally guaranteed, contingent income transfers. This calculation is hard to make precise, but it is likely that half of the \$42 billion annual benefit is, in fact, recoverable.

The forgoing discussion has focused on the negative side of these results. Under typical economic assumptions, it might be impossible to distribute these benefits to all members of the contracting generations. The answer depends on how we model individuals' lifetimes, on whether the size of future working generations can be approximated by knowing their numbers as children, and on whether age-specific, as well as generation-specific, transfers are possible within the pay-as-you-go system.

Notes

1. This section is based on Green 1977.
2. See, for example, Arrow 1965.

References

- Arrow, K.J. 1965. *Aspect of the Theory of Risk-Bearing*. Helsinki: Yrjo Jahnssonin Saatio.
- Boskin, M., and D. Puffert. 1987. "Social Security and the American Family." National Bureau of Economic Research Working Paper No. 2117.
- de Bartholome, C., and J. Brandts. 1986. "Social Security and Population Uncertainty: The Full Nature of Demographic Bias." Unpublished manuscript.
- Green, J. 1977. "Mitigating Demographic Risk through Social Insurance." National Bureau of Economic Research Working Paper No. 215.