NBER WORKING PAPER SERIES

INTERGENERATIONAL RISK SHARING

Roger H. Gordon

Hal R. Varian

Working Paper No. 1730

Intergenerational Risk Sharing

ABSTRACT

In this paper, we argue that in designing government debt and tax-transfer policies, it is important to consider their implications for the allocation of risk between generations. There is no reason to presume that the market or the family can allocate risk efficiently to future generations, implying that stochastic government policies have the potential to create first-order welfare improvements. The model provides a non-Keynsian justification for debt-finance of wars and recessions, as well as an added rationale for Social Security type tax-transfer schemes which aid unlucky generations, e.g., the Depression generation, at the expense of luckier generations.

Roger H. Gordon
Department of Economics
University of Michigan
Ann Arbor, MI 48109

Hal R. Varian
Department of Economics
University of Michigan
Ann Arbor, MI 48109

INTERGENERATIONAL RISK SHARING

by
Roger H. Gordon
and
Hal R. Varian
University of Michigan

Many years ago, Arrow and Debreu argued that economic efficiency requires that any risk in the economy be shared among individuals in such a way that each individual charges the same risk premium for an additional share in each lottery. If such risk sharing does not occur, then in principle contingent consumption can be reallocated so as to make everyone better off. For example, a number of authors have explored the implications of the fact that individuals cannot easily avoid bearing the risk in their own labor income. Given this observation, Varian[1980] and Eaton and Rosen[1980] have argued that a personal income tax may result in a more efficient allocation of risk in the economy.[1] Even when fluctuations in labor income are shared by all workers, Merton[1984], Fischer[1982], and Enders and Lapan[1982], have argued that when there is no market in a security corresponding to aggregate labor income then the government can still reallocate risk to nonworkers through a labor income tax in a Pareto-improving way.[2] Similarly, Aizenman[1981] argued that the government can beneficially offset domestic economic shocks through modification of the exchange rate when the international securities market fails to pool risks internationally.

In each of these cases, there exists the question of why there is no market in a particular lottery, and whether the reasons that prevented the market from pooling this risk appropriately may also prevent the government from raising utility by doing so. For example, if there is no private market allowing diversification in the risk in an individual's labor income because the moral hazard costs of doing so outweigh the benefits, then the government, also facing these moral hazard costs, can do no better than the market. In contrast, if no private market exists because of adverse selection reasons, then the government can avoid these problems by making participation in the risk pooling scheme compulsory. Since it is normally unclear why private market fails to exist in a particular lottery, it is also unclear whether there is any potential for fruitful government intervention.

Markets also do not appear able to pool lotteries faced by different non-overlapping generations. However, here the reason why risk-sharing markets fail to exist seems clear. Later generations cannot participate through the market in lotteries which occur before they are born because they are not alive ex ante to buy shares in these lotteries. An agent cannot profitably buy shares on behalf of later generations since there would be no legal

---

[1] Similarly, Buchanan[1975] argued that this risk sharing aspect of the personal tax may explain the strong political support for it.

[2] While there may be no explicit mechanism available to trade risk in labor income, firms may be able to reallocate risk between workers and capital owners through implicit labor contracts, an idea explored by Azariades[1975] and Bailey[1974].

1

mechanism to force these later generations to accept any losses, implying that the agent would have no incentive to pass on any gains. The problem is that later generations cannot precommit themselves to participate in a lottery even when they would gain in expected utility by doing so. By the time they can commit themselves, they know the outcome. These problems arise whenever two generations are not both active in the securities market before the outcome of a particular lottery is revealed, even if their lives do overlap at some point.

The government, however, could well have the power to precommit later generations to share in the outcome of earlier lotteries. If so, there is the potential for Pareto-improving government policies which share risk between different generations. In fact, some government policies do seem to have been used, whether consciously or otherwise, for just this purpose. For example, historically, debt has been issued to help finance unfavorable events, e.g., wars and recessions, and paid off only gradually over later generations. Diamond[1965] worked out in a very general setting just how debt issues result in a reallocation of wealth between generations. By the same argument, *stochastic* issues of debt result in a reallocation of *risk* between generations. Similarly, a plausible argument might be made that the initial role of the Social Security program was to aid the generations who lost both financial wealth and labor income during the Great Depression, at the expense of much later generations.

The objective of this paper is to explore the characteristics of an optimal government risk-sharing scheme, assuming that the government does have the power to precommit future generations. Since there are an arbitrarily large number of future generations to share in any particular lottery (e.g., today's recession), one might expect, by analogy with the diversification theorems in finance, that on efficiency grounds each generation ought to bear an arbitrarily small share in the outcome of any lottery.[3] This argument is formalized in section 1.

One problem with this argument is that the process by which random outcomes are shared with future generations is by adding to the capital stock when the outcome is favorable and consuming part of the capital stock when the outcome is unfavorable. Yet accounting explicitly for the economic effects of changing the capital stock makes the problem much more complicated. These and other complications are explored in section 2.

In these first two sections, we assume that the government does have the power to precommit future generations, and develop a normative model of government behavior. However, the government may not plausibly be able to precommit later generations, regardless of the ex post outcomes. While the risk-sharing policy raises the expected utility of future generations, based on the information available at the time of enactment, future generations at birth may well find themselves worse off due to the risk-sharing policy, given the observed outcomes on all past lotteries. In section 3, we develop two different positive models of government behavior to explore when a given risk-sharing scheme, present

---

[3] This intuition has been stated previously in Stiglitz[1983ab] and Gordon[1985].

2

initially, would end up being repealed at a later date when some future generation, seeing past outcomes, finds the implications of the policy unfavorable for them.

Throughout the above sections, we ignore the possibility that the family can engage in similar risk-sharing schemes between different generations, eliminating the need for government intervention. In fact, one might reinterpret our results as describing optimal family behavior rather than optimal government behavior. In section 4, we explore briefly in what ways risk-sharing through the family would face different problems than risk-sharing through the government.

1. Risk Sharing Between Generations: Base Case

We begin by exploring risk sharing in a very simple two period overlapping generations model. For simplicity, we assume there is no population growth and no technological change. Generations are therefore identical except for the period in which they are born. The generation born in period t, and its earnings, are denoted through use of a subscript t.

Members of each generation are assumed to work only while they are young and to consume only while they are old. We assume they work a fixed amount while young, earning a nonstochastic wage w. They save this labor income and earn a stochastic return, e_t, on it of mean zero, and variance s, allowing them to consume $w + e_t$ when they are old.[4] Their utility depends solely on their consumption, which is stochastic. We assume, for simplicity, that their ex ante utility can be expressed as a function of their expected income and the variance of this income, and denote their utility by $U(w) - V(s)$.[5]

The outcome of the stochastic event is revealed "between periods," so after the previous generation has died but before the next generation is born. For simplicity, we assume that each stochastic event is identically distributed and independent of all others. Because there are never two generations alive both before and after a stochastic event, there is no possibility of sharing risk between generations through the market. We assume, though, that risk is efficiently shared among members of each generation. Each individual therefore has expected utility of $U(w) - V(s)$. When it adds clarity, we include a subscript on the utility function indicating the generation number.

Since two generations are always alive simultaneously, however, the government can transfer income from one to the other, based on the outcome of past events, and raise the expected utility of every generation, expected as of the date of enactment of the program. For example, if the government transfers $e_{t-1}/2$ from the old to the young in each period t, then the young in period t receive net income of $w + e_{t-1}/2$. This income is saved, providing $w + e_{t-1}/2 + e_t$ the next period.[6] Part of this is paid as a transfer to the next

---

[4] This stochastic return represents the return on "the market". We assume that all idiosyncratic risks affecting subsets of the population have been diversified away.

[5] Most of our results generalize easily to an expected utility model, though exposition is simpler in a mean-variance setting.

[6] We assume that the same random income e_t is received independently of how much

generation, however, leaving $w + (e_{t-1} + e_t)/2$ for consumption. Expected utility of all but the initial generation t would therefore be $U(w) - V(\text{var}((e_{t-1} + e_t)/2)) = U(w) - V(s/2)$. Each lottery is shared between two generations, allowing a pooling of risks and an increase in expected utility. The first generation does yet better since it shares in only one lottery. Its expected utility equals $U(w) - V(\text{var}(e_1/2)) = U(w) - V(s/4)$.

This tax-transfer scheme can easily be redesigned so as to share each lottery between n generations. For example, the transfer from old to young in period t can be specified to be $\sum_{i=1}^{n-1} e_{t-i}/n$. If this policy has been in effect for at least $n-1$ generations, the utility of generation t, expected as of the date of enactment of the policy, can be expressed as $U(w) - V(\text{var}(\sum_{i=0}^{n-1} e_{t-i}/n)) = U(w) - V(s/n)$. Again the first $n-1$ generations would fare yet better. As risk is shared among more generations — as n becomes larger — expected utility of each generation increases as well. In fact, as n increases without bound, each generation's utility converges to $U(w)$ and the costs of bearing the collective risks drop to zero.

In this argument, we described the government policy as a tax-transfer scheme. It could equivalently have been described as a particular stochastic government debt policy. If each lottery is to be shared equally between n generations then the government upon seeing the outcome e_t can retire $[(n-1)/n]e_t$ dollars of debt, funded by a tax on the old (those who received e_t), then reissue e_t/n dollars of debt during each of the following $n-1$ periods, paying the proceeds of each debt issue to the old in that period. This policy results in the same redistribution of risk as the tax-transfer scheme described above.

Obviously, the above model is highly simplified. However, the conclusions are robust to many types of generalizations. The basic intuition being modelled is that each lottery ought to be shared relatively equally among all current and future individuals, and there are many future individuals relative to the current number of individuals. What is social risk at any one date is idiosyncratic risk when pooled with the independent lotteries of many later generations. Thus almost all risk should be passed forward to future generations. This conclusion remains even if we introduce an autocorrelation structure to the error terms, as long as no shock leaves permanent after effects. Similarly, we can allow individuals to consume throughout their life, have lotteries which occur at any point during their life, and have any type of risk-averse utility function, and still argue that idiosyncratic risks ought to be pooled. Allowing for income growth or population growth would only increase the incentive to share risk with the future.

2. Risk Sharing Between Generations: Complicating Factors

However, some of the simplifying assumptions in the above argument are critical. The mechanism by which random fluctuations in income are transmitted to later generations is through random fluctuations in the capital stock. In the above argument, these fluctuations in the capital stock had no effect on wage rates or interest rates, and we did not impose a nonnegativity constraint on the capital stock. Taking account of these complications results in a substantial quantitative change in the conclusions of the argument.

_____

is saved from the previous period.

One general way to respecify the problem, taking account of these complications would be as follow. Given the capital stock, k, saved by the older generation, and given the labor supply, h, of the younger generation, output equals $f(k,h)z$, where z is a random variable with mean one and some variance s, and where $f(.,.)$ satisfies the standard characteristics of a production function. In order to insure that the desired capital stock remains positive, we could assume that the production function satisfies the Inada conditions with respect to k, given h. Available resources for consumption and savings then equal output plus the remaining capital stock k. Tax or debt policy could be used as before to vary the capital stock stochastically to reallocate risk between generations.

This problem is related to those explored in the optimal stochastic growth literature by, for example, Merton[1975] and Brock-Mirman[1972]. Unfortunately, we have not been able to conclude anything of interest about an optimal stochastic tax policy in this general a setting. While the lack of an intergenerational risk-sharing market implies the possibility of Pareto improving trades accomplished through government intervention, little more can be said in general. Even the desired direction of transfer of risk is not obvious, for in the above formulation the wage rate and the return to capital will vary proportionately. The market therefore divides the lottery proportionately between the initial two generations. Since the second generation now saves a stochastic amount depending on this initial lottery, the wage rate of the next generation also depends on the outcome of this first lottery, and so on. There is nothing in the specification to indicate how this arbitrary division of risk between different generations differs from an efficient division.

In this section, we will try to shed some light on the factors affecting the nature of the optimal risk sharing policy by introducing one type of complication at a time and describing its implications for the optimal policy. In doing so, we make two simplifying assumptions about government policy. First, we assume that if the government is sharing each lottery between n generations, then it can choose any proportional amount a_i of the initial lottery to allocate to the generation born i periods later, subject to the constraint that $\sum_i a_i = 1$. For simplicity, the sharing rule is not allowed to vary over time, as a function of the state of the economy, or as a function of the outcome of the lottery.

The second simplifying assumption we make is that the objective of the government is to maximize the expected utility of the steady-state generations, expected as of the date when the policy is enacted. By steady-state generations we mean those generations who share in $n-1$ past lotteries in addition to keeping some share of their own lottery. In contrast, the first $n-1$ generations alive after the policy is enacted fare better than these steady-state generations since they keep the same fraction of their own lottery, but do not have to participate in as many past lotteries. Therefore, if a policy raises steady-state expected utility, it will raise the expected utility of earlier generations as well, and therefore be a Pareto improvement.

We realize that the definition of Pareto improvement in the context of differential information is somewhat controversial. We have adopted the view that a Pareto improvement means that each generation is expected to be better off given the information available at the time of enactment of the policy. An alternative definition of Pareto improvement, ex-

amined in Peled[1982], requires that each generation have higher expected utility at birth, conditional on the information available at that time. In the context of intergenerational redistribution, this would require that each generation have higher expected utility *given* the payments it owes to the previous generation. Peled shows that, under his definition, a competitive equilibrium is Pareto efficient, and thus no government intervention is warranted. When we claim a Pareto improvement, we mean by this that each generation gains by the policy change in the eyes of the decision-maker, evaluated using the information available at the time the policy is enacted.

Using the model of the economy of section 1 and the policy described there, a steady-state generation, born in period t, would end up with consumption equal to $w + \sum_{i=0}^{n-1} a_i e_{t-i}$ and expected utility of $U(w) - V(s \sum a_i^2)$. The policy which maximizes steady state utility is simply $a_i = 1/n$ for all i. Therefore the optimal policy, given the above assumptions is precisely the equal sharing policy examined previously in section 1. In the following sections, we will examine how this optimal policy changes as we change assumptions.

2.1 Positive Real Return to Savings

To begin with, we assume that income from capital equals $rk_t + e_t$, where r is greater than zero and nonstochastic, rather than simply e_t.[7] If generation t is still allotted the share a_i of the lottery which occurred i periods earlier, then its consumption in the second period would be $w + \sum_{i=0}^{n-1} a_i e_{t-i} (1+r)^i$, yielding expected utility of

$$U(w) - V[\sum_{i=0}^{n-1} a_i^2 (1+r)^{2i} s].$$ (1)

Maximizing steady state utility with respect to the a_i, subject to the constraint that $\sum_i a_i = 1$, we find at the optimum that

$$a_i = \frac{(1+r)^{-2i}}{\sum_{i=0}^{n-1} (1+r)^{-2i}}$$ (2)

and utility equals $U(w) - V(s/(\sum_{i=0}^{n-1}(1+r)^{-2i}))$. As before, utility increases as n increases, and in the limit utility converges to $U(w) - V(s(r^2 + 2r)/(1+r)^2)$, while a_i converges to $(r^2 + 2r)/(1+r)^{2(1+i)}$. Unlike in the base case, each generation still bears some risk and, in particular, passes forward to later generations only a fraction $1/(1+r)^2$ of its own lottery.

In the original model, when part of the outcome of a particular lottery was transferred to the next generation, the consumption of that generation changed by exactly the amount transferred. Now, however, when a lottery is passed to the next generation, its consumption changes by $(1+r)$ times the outcome of the lottery. Therefore, the variance of its consumption increases by $(1+r)^2$ times the variance of the outcome of the lottery. In the process of transferring risk to the next generation, it grows in size due to the earnings on the random addition to the capital stock. This change in the terms of trade between generations makes it more expensive to pass risk forward to later generations, and as a result less risk-sharing ought to occur.

---

[7] If we introduced population growth simultaneously, then r would simply be reinterpreted to equal the excess of the interest rate over the population growth rate.

2.2 Compounding of lotteries

For simplicity of notation, let us assume again that $r = 0$, but now assume that the stochastic shock to income is proportional to the capital stock, so that the income to capital equals $e_t k_t$ rather than simply e_t. When some share of the outcome of a lottery is passed forward to later generations, it earns a stochastic rate of return, so lotteries compound.

If each individual is still allotted the share a_i of the lottery occurring i periods earlier, then due to compounding of lotteries, the extra risk it would bear as a result of this lottery would be $a_i^2 s (1 + s)^i$. Steady state utility becomes $U(w) - V(s \sum_{i=0}^{n-1} a_i^2 (1 + s)^i)$. This expression is very close to that appearing in equation (1), and we can conclude directly that the optimal values of the a_i are

$$a_i = (1 + s)^{-i} / (\sum_{i=0}^{n-1} (1 + s)^{-i}),$$ (3)

and in the limit, as n increases, $a_i = s/(1+s)^{1+i}$. When risk is passed to a later generation, it grows in variability due to the random return it earns each period. As a result, it is more "expensive" than in our base case to pass risk on to later generations, and less ought to be done.

2.3 Diminishing returns to savings

When the policy proposed in section 1 is used to share risk among n generations, the capital stock saved by a steady state generation would equal $w + \sum_{j=1}^{n-1} ((n - j)/n) e_{t-j}$. There is nothing in the model which assures that this capital stock is positive. Moreover, as n gets larger, the stochastic process for the capital stock follows close to a random walk from one generation to the next, and a random walk has probability one of hitting any boundary (e.g., a requirement that the capital stock be nonnegative) in finite time. Our model therefore does not guarantee that net assets remain positive.

This is not such a problem with the model if the economy is open, and if the government can borrow from foreigners at an exogenous interest rate. The domestic capital stock can be large and positive even when the net asset position of individuals in the country is negative since foreigners can own large amounts of both corporate and government bonds. Even the world's capital stock is finite, however, so the logical constraint that the capital stock be nonnegative will eventually be violated even in an open economy.

One way to proceed is to impose a constraint that net assets must always be positive, and recalculate the optimal policy.[8] Introducing these constraints should lower the optimal degree of risk-sharing with the future, and also provide an added incentive to build up the capital stock.

---

[8] Such a model would be related to models of optimal savings as in Foley-Hellwig[1975], models of optimal commodity stockpiles, and also models of lifetime consumption and portfolio rules under uncertainty, as in Merton[1971].

The prime force keeping the capital stock positive, however, is presumably the large marginal product of capital when the capital stock is small. Assume, for example, that the income earned by capital equals $f(k_t) + e_t$, where $f(.)$ satisfies the Inada conditions. To avoid the complications added by a nonzero interest rate, assume that $f(w) = 0$.

In this setting, it is easy to show that some risk-sharing remains worthwhile. Consider, for example, a two generation tax-transfer scheme. Given the transfer policy, the utility of a generation in steady-state can be expressed as

$$U[w + Ef(w + ae_{t-1})] - V[(1-a)^2 s + a^2 s + \text{var}(f(w + ae_{t-1})) + Eae_{t-1}f(w + ae_{t-1})].$$

The first derivative of utility with respect to a equals

$$U'\frac{\partial Ef}{\partial a} - V'[\frac{\partial \text{var}(f)}{\partial a} - 2(1-2a)s + Ee_{t-1}f + aEe_{t-1}^2 f']. \tag{4}$$

Evaluated at $a = 0$, this derivative simplifies to $2sV'$, which is clearly positive. Therefore, some policy intervention raises steady-state utility.

The first-derivative described in equation (4) contains a number of extra terms which do not appear in the earlier models. The first term measures the drop in expected capital income when more risk is shared. Since the production function is concave, adding random fluctuations to the capital stock causes expected capital income to fall. In addition, the variance of consumption should increase faster than it did in the original model as risk-sharing increases. First, income from capital should become more variable as the capital stock becomes more variable.[9] The last two terms capture the increase in risk due to the positive correlation between income from capital and the size of the capital stock. Together, these four extra terms all reduce the gain from increasing a at every point, and imply that the optimal value of a is smaller here than in the original model.

2.4 Moral hazard effects of a tax-transfer scheme

In the original model, the tax-transfer scheme created no excess burden — individuals had an exogenous labor supply, saved everything from the first period, and consumed everything during the second period, regardless of the policy. In this section, we explore the possible moral hazard problems created by taxing random income.

What moral hazard problems are created depends on the specific form of taxation. Tax payments by any individual could be specified to be independent of that individual's actions, depending only on the outcome of the "market" as a whole. In this case, each individual's incentives are left unchanged by the transfer payment, even though risk is being shared with later generations. Of course, if individuals are heterogeneous, such a policy may not be desirable on distributional grounds.

If each individual is taxed based on his ownership share in the market lottery, then in general the tax does create a distortion cost, and these costs must be traded off with

---

[9] This is unambiguously true when a is small.

the efficiency gain from spreading risks across generations. For example, assume that each individual can invest in either of two assets, one riskless and the other risky. His income from savings will equal $i(w_a - k) + (r + e)k$, where w_a is his first-period income after including transfers, i is the riskless rate of return, k is the amount invested in the risky asset, r is the expected return on this asset, and e is the random component of the return. To maintain consistency with our original model, we simplify by setting $i = 0$. The government taxes each individual an amount equal to $\alpha(r + e - \beta)k$, for some tax parameters α and β, and transfers the proceeds in a lump-sum fashion to members of the next generation.

This tax scheme reduces to the two generation risk-sharing scheme described originally if $\beta = r$ and $\alpha = .5$. Given these parameters, we show that investors will not take account of the risk-bearing costs of that share of the risk going to the next generation when they make their investment decisions, resulting in overinvestment in the risky activity. As a result of this moral hazard cost, the optimal degree of risk-sharing would be smaller. However, the size of the moral hazard cost depends critically on the size of β. In particular, if the transfer to the next generation has an expected value just large enough to compensate them for bearing the extra risk, then we show that investment decisions are not distorted and the optimal degree of risk-sharing remains unchanged

Under the above tax-transfer scheme, consumption of a generation in steady-state[10] would equal $(w + \alpha(r + e_{-1} - \beta)k_{-1} - k) + (1 + r + e)k - \alpha(r + e - \beta)k$, where e_{-1} and k_{-1} refer to variables describing the previous generation. Expected utility, expected as of the date the policy is enacted, therefore equals

$$U[w + \alpha(r - \beta)Ek_1 + (r(1 - \alpha) + \alpha\beta)Ek] - V[\alpha^2 sEk^2_{-1} + (1 - \alpha)^2 sEk^2]. \qquad (5)$$

Differentiating steady-state utility with respect to α, setting the derivative to zero, and performing some simple manipulations yields

$$2\alpha sEk^2_{-1}V' = 2(1 - \alpha)sEk^2V' + (r - \beta)(Ek_{-1} - Ek)U' + [U'(r(1 - \alpha) + \alpha\beta)E\frac{\partial k}{\partial \alpha}$$

$$- V'2(1 - \alpha)^2 sEk\frac{\partial k}{\partial \alpha}] + [U'\alpha(r - \beta)E\frac{\partial k_{-1}}{\partial \alpha} - V'2s\alpha^2 Ek_{-1}\frac{\partial k_{-1}}{\partial \alpha}]. \quad (6)$$

Since the individuals in this generation choose k optimally, the third term on the right-hand side of equation (6) must equal zero. Also, at the date the policy is enacted, the expected steady-state capital stock is constant, so $Ek_{-1} = Ek$, implying that the second term on the right-hand side of the equation also equals zero. It therefore follows that the optimal α, denoted by α^*, must satisfy

$$\alpha^* = .5 + AU'(r - \beta)E\frac{\partial k_{-1}}{\partial \alpha} - AV'2s\alpha^* Ek_{-1}\frac{\partial k_{-1}}{\partial \alpha}, \qquad (7)$$

---

[10] When behavior changes in response to the tax, k will evolve over time after the tax-transfer scheme is enacted. In steady-state, as now defined, the probability distribution of k remains constant.

9

where $A = \alpha^*/(4sV'Ek^2)$.

If $\beta = r$, as was implicitly assumed in the original model, then the second term in equation (7) equals zero. In addition, since the government now absorbs a fraction of the random component of the return on the risky activity without changing the mean return, it is easy to show that $(\partial k_1/\partial\alpha) > 0$. Each generation chooses to undertake too much of the risky activity, since it ignores the risk-bearing costs imposed on the next generation, and we see from equation (7) that the optimal value of α is now smaller than .5.

This moral hazard cost disappears, however, if the last two terms in equation (7) sum to zero, which occurs if any extra risk transferred to the next generation is offset by a suitable risk premium. For example, if $\beta = 0$, we find comparing the third and fourth terms on the right-hand side of equation (6) that if $\alpha = .5$, then the fourth term must equal zero as well. If the fourth term equals zero then equation (7) implies that the optimal α does in fact equal .5. With $\beta = 0$ and $\alpha = .5$, each generation is in effect a fifty percent partner in the risky investment undertaken by the previous generation. Given the symmetric positions of the two generations, when one generation is indifferent to further investment, so is the other. Investment decisions are no longer distorted, and there are no moral hazard costs generated by the tax-transfer scheme.

In general, eliminating any distortion in the risk-sharing scheme requires that future generations be just indifferent to the change in transfers that results when an investor considers changing his investment in a risky activity. Except in this situation, moral hazard problems do complicate the design of a tax-transfer scheme.[11] Moral hazard problems do not necessarily lower the amount of risk shared with the future, however. For example, if β were less than zero, then it is straight-forward to show that the optimal α^* is greater than .5.

3. Time Consistency of the Policy

In analyzing the government policy so far, we have assumed that what is optimal ex ante will in fact be done ex post. However, at each future date, both generations know the outcome of past lotteries. Even if each generation gains ex ante under the risk-sharing policy, a generation may well find, after seeing the outcome of past lotteries, that it loses from the policy and would wish to repeal the program. Whether we would expect it to succeed in repealing the policy, however, depends on our model of political decision-making. In this section, we examine the political stability of a risk-sharing scheme in two alternative models of political decision-making.

In both models, we impose two restrictive assumptions. First, we assume that voters can repeal an existing risk-sharing policy, but cannot otherwise change its design. Without this assumption, each generation would try to obtain nonstochastic transfers at the expense of other generations. In addition, we assume that if the risk-sharing policy is repealed in a particular period, then it cannot be reenacted in the next period. Otherwise, each

---

[11] If the amount transferred to an individual depends on his behavior, then further moral hazard complications arise.

policy they would follow. In general, however, the family would face a number of difficulties not faced by the government.

If parents do expect to be able to share their lotteries with future generations within their family, then they would wish to maintain close to their normal consumption level even during a very unfavorable event, e.g., a depression. However, they may not at that point have the financial assets to fund this consumption level. Under the government risk-sharing policy, the government would issue debt and use it to finance transfer payments. Under the analagous family policy, the parents would try to borrow funds to finance their consumption. However, they would lack collateral for the loan, and could not legally commit their children to pay it back. In contrast, the government would not need explicit collateral and can legally bind future generations whenever it issues debt.

As we saw in the previous section, a political coalition might well develop under certain circumstances to repeal a government risk-sharing policy. Within a family risk-sharing policy, each generation would face similar incentives to "repeal" the policy. However, within the family, the paying generation could repeal the policy on its own regardless of the preferences of the receiving generation. Repealing a family risk-sharing policy is much easier than repealing a government policy.

When would a generation within the family wish to "repeal" the risk-sharing policy? When considering the preferences of a generation under a government risk-sharing policy, we assumed that there was a clear link between its willingness to participate in the policy when it is young and its likelihood of participating in the policy when it is old. In particular, we assumed that if the policy were repealed now, it would not exist during the next period. Within the family, however, there is no contractual link between transfers now and transfers next period. Children may well hope that their children will help them out, independently of whether they help their parents out. Without this link between transfers now and transfers next period, every generation would face an incentive to renege. Altruism may overcome this, but without sufficient altruism the family risk-sharing policy is likely to break dc•·n.

5. Conclusions

In this paper, we have argued that in designing government debt and tax-transfer policies, it is important to consider their implications for the allocation of risk between generations. There is no reason to presume that the market or the family can allocate risk efficiently to future generations, implying that stochastic government policies have the potential to create first-order welfare improvements. The model provides a non-Keynsian justification for debt-finance of wars and recessions, as well as an added rationale for Social Security type tax-transfer schemes which aid unlucky generations, e.g., the Depression generation, at the expense of luckier generations.

It is premature, though, at this point to draw any conclusions concerning the degree to which government debt or tax-transfer policy ought to be modified in light of risk-sharing considerations. Given the theoretical difficulties of characterizing the optimal policy in more general settings, the answer to this question must await more detailed theoretical

work as well as results from computer simulations.

generation would try to avoid sharing in past lotteries, yet attempt to get future generations to share in its own lottery. This restriction seems plausible given the stability of existing government policies.

In our first model of political decision-making, we assume that the younger generation is slightly more numerous and that the median voter, who would be young, controls the government.[12] Assume in addition that each generation's lottery is bounded below by $-b$, and that the proposed policy involves a transfer from children to parents which is a strictly decreasing function of the outcome of each of the previous $n - 1$ lotteries. Assume as well that the policy involves a net payment from children to parents when the worst outcome occurs during each of the previous $n - 1$ lotteries. But if all these worst outcomes do occur, the children will surely vote for repeal. They must make the maximum possible payment to their parents for certain, and they can receive back the same amount from their own children only if they also suffer the worst outcome on their lottery. With any other outcome of their own lottery, they lose on net from the policy. Assume then that there is some maximum payment that children would willingly make to their parents without voting for repeal. But if children are asked to make this payment, by the same argument they will refuse. At best, they can just recoup their losses from their own children, and most of the time they lose on net. Therefore children would repeal the policy whenever they are asked to make a payment to their parents, and would even repeal it when the transfer their parents make to them is sufficiently small.[13]

Once the policy has been repealed, however, the next generation would have no interest in restarting the policy, since regardless of its own luck, its children would never pay them anything, whereas the policy would still involve its paying money to the next generation under some outcomes. They would be better off with no policy in effect. Therefore, even if the policy were in effect initially, once children are asked to make a payment to their parents, the policy would be repealed and never reestablished.

One obvious mechanism to handle this likelihood of repeal is to require a sufficiently large plurality for repeal. Another would be to build into the policy a large enough penalty on any generation that chooses to repeal to eliminate any incentive to do so. This penalty could involve either a tax, the proceeds from which could be used to retire government debt, or merely the administrative and social disruption that would ensue if the government were to default on its legal obligations (as would happen if past transfers were funded by government debt rather than by a tax-transfer scheme, such as Social Security). To guarantee that repeal would never occur, however, this penalty must be large enough so that even if the worst outcome has occurred during the previous $n - 1$ generations the next

---

[12] Population growth affects the optimal policy in a continuous fashion, but affects the median voter discontinuously. We ignore the effects of slight population growth on the optimal policy.

[13] This argument breaks down if population growth is important, for then much more is received from the next generation than is paid to the previous generation. But with population growth and a zero interest rate, the economy is beyond the golden rule implying various possible ways to increase everyone's utility.

generation would rather follow through on the policy, receiving in partial compensation the right to share its own lottery with future generations, rather than face the penalty.

One unattractive feature of the median voter model is that it implies that the Social Security program should have no chance of survival politically, since it always requires that children pay parents. Yet historically Social Security has been a remarkably stable program. As one simple alternative to the median voter model of political decision-making, let us assume that a policy would be repealed whenever those in favor of repeal are willing to pay more towards this effort than those against repeal are willing to pay to avoid it.[14] The government, for example, could be presumed to support that policy receiving the most support from lobbiests. Given this model of government decision-making, the analysis of when repeal occurs is very simple. At any date, the optimal risk-sharing policy involves a nonstochastic transfer between old and young plus a promise of a stochastic transfer to the current young during the next period. Ignoring the next period's transfer, both generations have exactly the same income at stake, and their political influences would exactly counterbalance by assumption. If the presence of the stochastic transfer during the next period raises the expected utility of the current young, conditional on knowledge of the outcome of past lotteries, then they would be more in favor of maintaining the policy, or at least less inclined to oppose it. Since preferences were otherwise exactly counterbalanced, we conclude that the policy will remain in force whenever the young gain by the presence of the stochastic transfer during the next period.

When do the young gain from the next period's transfer? They certainly gain if the policy always involves a (stochastic) payment from young to old, as under Social Security. If the policy were always to require payments from parents to children, then it would certainly be repealed. In general repeal would occur whenever, given knowledge of the outcome of past lotteries, the expected payment from parents to children during the next period is large enough to offset any gain from sharing risk with future generations. Thus, in sharp contrast to the results of the previous model, we now find that a tax-transfer program like Social Security has a most stable design, and according to our model would never be repealed.

Under this model, if a risk-sharing policy does not yet exist, any generation would be very much in favor of enacting it, since the first generation receives a transfer from its children while making none to its parents. Why should we expect the political process to enact a risk-sharing program, however, rather than simply to transfer wealth from future to present generations? We have not attempted to address this more basic question.

4. The Family as an Alternative Risk-Sharing Institution

Why cannot the family provide the same type of risk sharing between generations of the family that the above government policies would provide? If parents are altruistic and would always wish to leave a positive bequest, regardless of events, then they could easily provide the same risk-sharing. The above analysis in effect characterizes the risk-sharing

---

[14] See Wellisz and Wilson[1984] for an extended application of this model.

REFERENCES

Aizenman, Joshua, "The Use of the Balance of Payments as a Shock Absorber in Fixed Rate and Managed Float Systems," *Journal of International Economics* 11(1981), 479-486.

Azariades, C., "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy* 83(Oct., 1975), 1183-1202.

Baily, M. N., "Wages and Employment Under Uncertain Demand," *Review of Economic Studies*, 41(Jan., 1974), 37-50.

Battacharya, S., "Aspects of Monetary and Banking Theory and Moral Hazard," *Journal of Finance*, 37(May, 1982), 371-384.

Brock, W. and L. Mirman, "The Stochastic Modified Golden Rule in a One Sector Model of Economic Growth with Uncertain Technology," *Journal of Economic Theory* 4(June, 1972), 479-513.

Buchanan, J., "Taxation in Fiscal Exchange," *Journal of Public Economics* 6(July, 1976), 17-30.

Diamond, P., "National Debt in a Neoclassical Growth Model," *American Economic Review*, 55(Dec., 1965), 1126-1150.

Eaton, J. and H. Rosen, "Taxation, Human Capital and Uncertainty," *American Economic Review*, 70(Sept., 1980), 705-715.

Enders, W. and H. E. Lapan, "Social Security Taxation and Intergenerational Risk Sharing," *International Economic Review*, 23(Oct., 1982), 647-658.

Fischer, S., "Welfare Aspects of Government Issue of Indexed Bonds," N.B.E.R. Working Paper No. 874, April, 1982.

Foley, D. and M. Hellwig, "Asset Management with Trading Uncertainty," *Review of Economic Studies*, 42(July, 1975), 327-346.

Gordon, R., "Taxation of Corporate Capital Income: Tax Revenues vs. Tax Distortions," *Quarterly Journal of Economics* 100(Feb., 1985), 1-27.

Merton, R. C., "An Asymptotic Theory of Growth Under Uncertainty," *Review of Economic Studies* 42(1975), 375-393.

——, "On the Role of Social Security as a Means for Efficient Risk-Bearing in an

Economy where Human Capital is Not Tradeable," in Zvi Bodie and John Shoven, eds., *Financial Aspects of the United States Pension System*, University of Chicago Press, 1984.

———, "Optimal Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory* 3(1971), 373-413.

Peled, D., "Informational Diversity over Time and the Optimality of Monetary Equilibria," *Journal of Economic Theory* 28(1982), 255-274.

Stiglitz, J., "On the Relevance or Irrelevance of Public Financial Policy," N.B.E.R. Working Paper No. 1057, Jan. 1983.

———, "On the Relevance or Irrelevance of Public Financial Policy: Indexation, Price Rigidities, and Optimal Monetary Policies," in *Inflation, Debt and Indexation*, ed. by Rudiger Dornbusch and Mario Henrique Simonsen, Cambridge: M.I.T. Press, 1983.

Varian, H., "Redistributive Taxation as Social Insurance," *Journal of Public Economics*, 14(Aug., 1980), 49-68.

Wellisz, S. and J. D. Wilson, "A Theory of Tariff Formation," mimeo, 1984.