

Preliminary

**Intergenerational Risk Sharing in the Spirit of Arrow, Debreu,
and Rawls, with Applications to Social Security Design**

Laurence Ball

Johns Hopkins University

N. Gregory Mankiw

Harvard University

March 2001

Abstract

This paper examines the optimal allocation of risk in an overlapping-generations economy. It compares the allocation of risk the economy reaches naturally to the allocation that would be reached if generations behind a Rawlsian "veil of ignorance" could share risk with one another through complete Arrow-Debreu contingent-claims markets. The paper then examines how the government might implement optimal intergenerational risk sharing with a social security system. One conclusion is that the system must either hold equity claims to capital or negatively index benefits to equity returns.

I. Introduction

Do market economies allocate risk efficiently? If not, what government policies can improve the allocation of risk? These are classic questions of economic theory. One celebrated answer comes from the Arrow-Debreu theory of general equilibrium. This theory teaches that under certain conditions--in particular, if contingent-claims markets are complete--the allocation of risk will be Pareto efficient. In other words, with complete markets, society can let the invisible hand allocate risk.

This paper explores a deviation from Arrow-Debreu theory that arises from a simple fact that not everyone is born at the beginning of time. In an overlapping-generations economy, markets must be incomplete, because a person cannot engage in risk-sharing trades with those who are not yet born. The risks associated with holding capital assets, for instance, can be shared with others alive at the same time, but they cannot be shared with future generations. As a result, the allocation of risk need not be efficient, and government policy may be able to make Pareto improvements.

The suboptimality of risk allocation in stochastic overlapping-generations models has been discussed in several recent papers, including Bohn (1998), Shiller (1999), Rangel and Zeckhauser (1998), and Smetters (2000). We approach this issue by considering a simple thought experiment. Imagine that all

generations ever to be born were here today and able to trade in complete contingent-claims markets. How would the allocation of risk in this complete-markets setting differ from the one the economy reaches without these prenatal risk-sharing trades?

This approach builds on two traditions. The first is the Arrow-Debreu theory of general equilibrium. In essence, our thought experiment opens up all markets that are assumed to exist in Arrow-Debreu theory but, in fact, cannot exist in an overlapping-generations economy. The second tradition is the Rawlsian approach to social justice. Our thought experiment envisions a hypothetical time period when all generations are alive in an "original position" behind a "veil of ignorance." In Rawls's (1971) work on social insurance, the ignorance concerns cross-sectional uncertainty about one's station in life. Here, the ignorance concerns time-series uncertainty about whether one is born into a lucky or unlucky generation.

This theoretical investigation is motivated by practical issues of public policy. The government influences the allocation of risk among generations in many ways, most notably through the social security system. A benevolent policymaker might try to use these instruments to achieve the allocation of risk that the invisible hand would reach if it could. That is, the policymaker might try to implement the outcome that people would achieve on their own if, as in our thought experiment, they were able to fully trade risks. Our goal, therefore, is not only to examine how

different the world would be with complete markets but also to discuss how, without such markets, government policy might substitute for them. This analysis sheds light, for instance, on how the social security trust fund should be invested and how taxes and benefits should respond to macroeconomic shocks.

We proceed as follows. Section II presents a stochastic overlapping-generations model that we use in our main analysis. To keep things simple, we assume a single source of risk: uncertainty about the return on capital. We begin by describing the equilibrium in which people can trade only with others alive at the same time, so each generation bears the entire risk realized during its lifetime. We call this the Hobbesian equilibrium, because it is the equilibrium that nature gives us (and because a person's involvement in the market economy is "nasty, brutish, and short.")

In Section III, we introduce the central thought experiment of the paper. Maintaining the overlapping-generations framework of Section II, we posit the existence of complete contingent claims markets--markets for the consumption good in each period in every possible history. The individuals who will make up all generations participate in these markets in an "original position" that exists before the beginning of time. We call the allocation of consumption determined in the original position the Rawlsian equilibrium.

Sections IV and V solve for the Rawlsian equilibrium in the model. In general, one cannot obtain an analytic solution, so we

simplify the problem in alternative ways. In Section IV we consider a special case in which only a single generation faces uncertainty in capital returns, and examine how this risk is shared with other generations. In Section V we consider the more general case in which every generation faces uncertainty. We derive an approximate solution that is valid when the shocks are small.

The solution we find in Section V takes a simple and intuitive form. We find that capital-return risk in each period is shared equally among the generations alive during that period and all subsequent generations. In contrast to the Hobbesian equilibrium, where consumption is serially uncorrelated from generation to generation, consumption in the Rawlsian equilibrium follows a random walk.

We also draw a connection in Section V between our problem and the well-known Ramsey social planning problem. The Rawlsian outcome we examine is determined by decentralized trading among an infinite number of generations. Nonetheless, the resulting path of consumption is the same as would be chosen by a Ramsey social planner with a particular discount rate on future generations' utility. Thus, the Rawlsian approach to generational risk offers a rationale for studying Ramsey optima.

We next move closer to issues of policy. Because it is not yet feasible to transport people back in time to an original position, free markets are not a practical way to share intergenerational risks. In Section VI, we ask whether government

policies can substitute for the missing markets and ensure the Rawlsian allocation. We find a simple policy that does so: a fully-funded social security system in which the system's trust fund holds equity. In this system, benefits are permanently adjusted in response to shocks to equity returns to keep the system solvent.

There is, however, more than one way for policy to achieve any given allocation of risk. Policymakers can also implement the Rawlsian equilibrium if the social security trust fund holds safe debt. Yet in this system, benefits must be adjusted in what, at first glance, may seem a surprising way: they must be negatively indexed to equity returns.

Section VII sketches extensions of our basic model, including one in which both wage growth and capital returns are uncertain. Here, the Rawlsian equilibrium shares both kinds of uncertainty across generations, and this allocation can again be implemented with an equity-based social security system.

Section VIII concludes.

II. The Model and the Hobbesian Equilibrium

This section describes the basic version of our overlapping-generations model, which we keep as simple as possible to build intuition. A fixed number of people is born each period, and everyone lives for two periods. When young, a person supplies labor inelastically and receives a fixed wage, which we normalize

to one. Each person consumes only when he is old. Thus a worker saves his entire wage when he is young and consumes the return on this saving, including the principal, when he is old. There are no bequests. We assume log utility over consumption; thus lifetime utility for an individual born at period t is

$$(1) \quad U_t = \log(c_{t+1}),$$

where c_{t+1} is the individual's consumption in period $t+1$, when he is old.

Our assumption that only the old consume simplifies the model by eliminating the intertemporal consumption decision of the young, which is not essential to the issue of intergenerational risk-sharing. In Section VII, we discuss a more general model in which individuals consume in both periods of life.

We let R_t denote the gross return on savings between periods $t-1$ and t . Because the generation born at t receives a wage of one and saves it all, its consumption at $t+1$ is R_{t+1} . We take the return R_t to be an exogenous random variable. These assumptions about factor returns could be justified by positing a linear technology for output: $Y = L + RK$.

We assume that R_t is distributed independently over time and has a two-point distribution. Let ρ be the average value of the interest rate. R_t equals $\rho+x_t$ with probability one half and $\rho-x_t$ with probability one half. x_t measures the degree of dispersion in the capital return in period t . It is natural to focus on the case in which x_t is the same for all t (and we will do so below), but it

will prove useful to have in hand the general case in which x_t varies across periods. We assume that the lowest possible return, $\rho - x_t$, is greater than one; this assures that the model satisfies Abel et. al's (1989) condition for dynamic efficiency in Diamond's (1965) sense.

As a benchmark, we begin by considering the equilibrium of this model without any intergenerational risk sharing. This equilibrium is assumed in most previous work on stochastic overlapping-generations models. It is based on the realistic assumption that people can trade only with others who are alive at the same time. We also make the standard assumption that the young enter the market after the current return on savings is realized; thus there is no remaining uncertainty within a period for the old and young to share. As noted above, we call this outcome the Hobbesian equilibrium.

Given the environment just described, the Hobbesian equilibrium is trivial to derive. The generation born at t saves its wage and consumes its wealth when old. Thus, $c_{t+1} = R_{t+1}$ and $U_t = \log(R_{t+1})$. Because R_{t+1} is random and uncorrelated over time, each generation bears all of a single idiosyncratic risk--the return risk in the period it happens to be old. Note that consumption is independently distributed from generation to generation.

This Hobbesian equilibrium is clearly inefficient. There would be gains if the old at period $t+1$ could share the risk it faces with the generations born at $t+1$ and later. However, by the

time these generations are born and ready to participate in markets, the outcome for period $t+1$ is already realized, and private improvements in risk allocation are no longer possible.

III. The Rawlsian Equilibrium

We now consider a hypothetical world with markets for intergenerational risk-sharing. We assume that all generations are placed in an "original position" that exists before period one, when the first generation is born. In this original position, each person knows when he will be born, but he does not know the future evolution of the economy; in particular, he does not know whether his generation will be lucky or unlucky in its realization of capital returns. In the original position, everyone can share the risks they face by participating in contingent-claims markets.

We use the following terminology. The "state" in period t is the realization of R in that period. Because there are two possible realizations of R , there are two possible states in each period. A "history" of the economy through period t is a sequence of states for periods $1, \dots, t$. There are 2^t possible histories of the economy through period t , which we index by $h=1, \dots, 2^t$.

We assume that the markets in the original position are complete in the sense that there is a market for the consumption good in each period and each possible history of the economy through that period. We index consumption at period t in history h by th . For each period t , there are 2^t markets for history-

contingent consumption.

A person born in period t receives a wage of one in all histories of the economy through t . Thus, his endowment is one unit of good th for all $h=1, \dots, 2^t$. He can sell part of this endowment and use the proceeds to buy conditional consumption goods at $t+1$. He can also save and thereby transform goods dated at t into goods dated at $t+1$. For each unit of good th that he saves, he receives units of goods $(t+1)h'$ for each history h' through $t+1$ that is a continuation of history h . There are two such histories: the saver receives $\rho+x_{t+1}$ of the good in one history and $\rho-x_{t+1}$ in the other. Given these possibilities, an agent chooses a basket of goods dated at $t+1$ to maximize the expected value of utility, $\log(c_{t+1})$.

A Walrasian auctioneer finds the price that equilibrates supply and demand in each market for history-contingent consumption. We are interested in the set of equilibrium prices and the resulting allocation of consumption in all possible histories of the economy.

Obviously, this is a hard problem, and we have been unable to find a general analytical solution. Therefore, we simplify the problem in two different ways in the next two sections. In Section IV, we assume that capital returns are uncertain in only a single period; thus there is only one shock. In Section V, we return to the general case of many shocks but solve the model using a first-order approximation. Thus, we derive a solution that is valid when

the shocks to capital returns are small. Each of these two special cases yields its own insights into the nature of the Rawlsian equilibrium.

IV. The Case of a Single Shock

The original position we have described includes many markets for sharing risks. In this section, we consider a special case in which only one generation faces uncertainty. In this example, the uncertainty concerns the capital return in period j for the generation born at $j-1$. In the notation introduced above, $x_j > 0$ and $x_t = 0$ for all $t \neq j$. This example helps develop intuition about the model, and it is a building block for the more general analysis below.

A. Solution for the Rawlsian Equilibrium

In this example, the possible histories of the economy collapse to a simple set. In period j , there are two possible states of the economy: the good state (G) in which $R_j = \rho + x_j$ and the bad state (B) in which $R_j = \rho - x_j$. $R = \rho$ in all other periods. There is only one possible history of the economy through t for $t < j$, and two possible histories through $t \geq j$: the history with state G at j and the history with state B at j . Thus there is only one market for consumption in each period $t < j$, and two markets for each period $t \geq j$. We index the goods in the various markets by t for $t < j$ and by t_i , $i=G$ and $i=B$, for $t \geq j$.

The Appendix describes in detail the solution for the Rawlsian

equilibrium. Here, we sketch the approach and results.

The starting point is the budget constraints and objective functions of individuals in the original position. Using these, we solve for history-contingent consumption demand as a function of the relative prices of contingent consumption goods and then for the equilibrium relative prices. In equilibrium, there is no trade in goods dated before j , when the shock occurs. The only motive for trade is to share the risk from the shock, and goods before j cannot be contingent on the shock. The key prices are those for contingent goods dated j and later; we denote these prices by P_{tG} and P_{tB} . The Appendix shows that equilibrium relative prices satisfy

$$(2) \quad P_{ti}/P_{(t+1)i} = \rho \quad \text{for } i = G, B;$$

$$(3) \quad P_{tB}/P_{tG} = Q \equiv [\rho^2 + (\rho-1)x_j]/[\rho^2 - (\rho-1)x_j].$$

These equations fully describe the path of equilibrium relative prices.

These equilibrium conditions are simple to interpret. Condition (2) concerns the prices of the good in different periods but the same realization of history. It is a no-arbitrage condition. In this economy, people trade consumption between t and $t+1$ both by saving and by participating in markets. In equilibrium, the two activities must yield the same return.

Condition (3) gives the relative price of consumption in the good and bad histories. This price is the same for all periods $t \geq j$; this follows from a no-arbitrage condition and the fact that

agents can trade across periods at a fixed rate. The key result is that $Q > 1$: it costs more than a unit of consumption in the good history to buy a unit in the bad history. This is necessary to induce agents to demand greater consumption in the good history, when more resources for consumption are available.

The Appendix also derives equilibrium consumption of each generation. For generations born before $j-1$, consumption in the Rawlsian equilibrium is the same as in the Hobbesian equilibrium. For the generation born at $j-1$ (the one that experiences the shock), consumption when old is given by

$$(4) \quad c_{jG} = [(1+Q)\rho + (1-Q)x_j]/2$$

$$c_{jB} = [(1+Q)\rho + (1-Q)x_j]/2Q .$$

For all generations born at j and later, consumption in the Rawlsian equilibrium is

$$(5) \quad c_{tG} = \rho(1+Q)/2;$$

$$c_{tB} = \rho(1+Q)/2Q, \quad t \geq j+1.$$

Equations (2) through (5) fully describe the Rawlsian equilibrium.

B. Discussion of the Rawlsian Equilibrium

The solution we have just described has two notable properties. First, the ratio of consumption in the good and bad histories is equal to $Q > 1$ for all generations born at $j-1$ and later. All these generations--those who are old when the shock occurs and those who come later--suffer the same proportional loss in consumption from a bad shock. In other words, the risk from the shock is spread equally across generations.

This contrasts sharply with the Hobbesian equilibrium. In that equilibrium, the return risk in period j affects only the old in that period. The ratio of consumption by the old in period j in the two histories is $(\rho+x_j)/(\rho-x_j)$, which is greater than Q . Thus, this generation reduces its risk by moving from the Hobbesian equilibrium to the Rawlsian equilibrium, where it can share risk with future generations.

The second notable result concerns average consumption in the Hobbesian and Rawlsian equilibria. For generations born at j and later, average consumption over the good and bad histories is $[(1+Q)^2/4Q]\rho$. This exceeds ρ , which is these generations' consumption in the Hobbesian equilibrium. Thus, average consumption is higher in the Rawlsian equilibrium than in the Hobbesian equilibrium for all these generations. A bit more algebra shows that for the generation born at $j-1$, average consumption is lower in the Rawlsian equilibrium. Of course, for all generations, utility must be higher in the Rawlsian equilibrium, for the Hobbesian allocation is still feasible.

These results have a simple interpretation. In the Hobbesian equilibrium, the generation born at $j-1$ is uniquely disadvantaged: it is the only generation facing return uncertainty. In the original position, it reduces this uncertainty through the contingent-claims markets. In essence, it buys insurance from later generations. But later generations are willing to sell insurance only if they are compensated for taking on the risk.

This compensation is reflected in a value of Q greater than one. As a result, later generations obtain more consumption in the good history than they give up in the bad history.

V. The Case of Many Small Shocks

Having explored the special case of a single shock, we now examine a more general case in which there are shocks in every period: x_t is positive for all t . In the original position, there are now 2^t markets for history-contingent consumption in period t . These markets yield rich opportunities for sharing risks among generations.

A. Solution for the Rawlsian Equilibrium

We are interested in solving for equilibrium consumption in all possible histories, of which there are 2^t through period t . We simplify this hard problem by using a first-order approximation that is valid as long as the shocks, x_t , are small. That is, we derive the Rawlsian equilibrium when there are small fluctuations in capital returns.

The advantage of using a first-order approximation is that it eliminates any possible interaction among the shocks in different periods. (This is shown formally in the Appendix.) Thus, we can use the results in the previous section to show the effect of any individual shock, and we can find the effect of a series of shocks by summing the effects of the shocks.

Consider, then, a shock in period j . Equations (3)-(5) show

consumption for period j and after. Substituting the expression for Q in equation (3) into equations (4)-(5) yields consumption in terms of the size of the shock, x_j , and the average return ρ . Taking a first-order approximation in x_j around $x_j=0$ yields

$$(6) \quad \begin{aligned} c_{tG} &= \rho + [(\rho-1)/\rho]x_j; \\ c_{tB} &= \rho - [(\rho-1)/\rho]x_j. \end{aligned}$$

(If you really want to see the details, go to the Appendix.)

According to equation (6), the shock to the capital return causes consumption to rise or fall by a fraction $(\rho-1)/\rho$ of the shock for each generation born at $j-1$ and later. Note there is no distinction here between the generation born at $j-1$, who lives through the shock, and later generations. They share the risk equally and in an actuarially fair way, so that all generations have the same average consumption. The previous result that later generations have higher average consumption no longer holds, because the relative price of consumption in the good and bad histories approaches one as the shock becomes small. That is, the compensation future generations demand to take on risk is second-order, so it vanishes as the shock becomes small.

While equation (6) shows how consumption responds to a single, small shock, the result for a series of small shocks is found by summing the effects of each shock. To express equilibrium in a particular history, we let θ_t be an indicator variable equal to one in the good state and minus one in the bad state. The history of the economy through t is given by the sequence $\{\theta_1, \theta_2, \dots, \theta_t\}$.

In any history, consumption in the Rawlsian equilibrium is given by

$$(7) \quad c_t = \rho + \sum_{j=1}^t \theta_j [(\rho-1)/\rho] x_j.$$

If the shocks are all the same size ($x_j=x$ for all j), then this expression reduces to

$$(8) \quad c_t = \rho + (N_t^G - N_t^B) [(\rho-1)/\rho] x,$$

where N_t^G is the number of periods through t with good realizations of the shock and N_t^B is the number with bad realizations. A generation's consumption is raised by a fixed amount for every good shock in the past and reduced by the same amount for every bad shock.

Note that the last equation implies

$$(9) \quad c_t - c_{t-1} = \theta_t [(\rho-1)/\rho] x.$$

In each period, the change in consumption is proportional to the current shock. Thus, even though consumption in the Hobbesian equilibrium was serially uncorrelated, consumption in the Rawlsian equilibrium follows a random walk. The reason is that full risk-sharing causes each shock to be spread equally over current and future generations. Rather than a shock affecting only the generation living through it, it affects later generations as well. Intergenerational risk sharing makes the impact of a shock both smaller and more persistent.

B. Rawls Meets Ramsey

Our random-walk result for Rawlsian consumption may seem familiar: it resembles Hall's (1978) result for the optimal plan

for a single, infinite-horizon Ramsey consumer under uncertainty.

The resemblance is not a coincidence. Any Pareto-efficient equilibrium solves a social planner's problem for some set of weights on the welfare of different agents. In our model, the Rawlsian equilibrium solves a planner's problem for maximizing a weighted average of the utility of different generations. For this planner's problem to become the Ramsey model, the only missing ingredient is for the weights to decline exponentially over time.

One can show the weights do in fact decline exponentially with the discount factor $\beta=1/\rho$. The following is a sketch of the proof. Consider an infinite-horizon social planning problem with log utility, a gross interest rate ρ , and an arbitrary discount factor $\beta(t)$ for discounting utility between periods t and $t+1$. Under certainty, consumption grows between t and $t+1$ at a rate $\beta(t)\rho$. With small shocks, certainty-equivalence holds, and expected consumption growth equals $\beta(t)\rho$. In our solution for the Rawlsian equilibrium, however, consumption is expected to remain constant. Thus, the discount factor must be the constant $\beta=1/\rho$. This establishes that the social planning problem associated with the Rawlsian equilibrium is the Ramsey problem with this particular discount factor.

This correspondence between the Rawlsian equilibrium and the Ramsey model is noteworthy. Economists often use Ramsey problems when studying the allocation of consumption over time. One justification for this approach is Barro's (1974) model, where the

discount factor measures the extent of altruism among generations.

The Rawlsian equilibrium, however, suggests a different rationale for studying the Ramsey problem and a different interpretation of the discount factor. In our environment, to replicate the Rawlsian equilibrium, the social planner has to distribute risk optimally but not pursue deterministic transfers across generations. If the discount factor β were not equal to $1/\rho$, the planner would want to move resources from earlier to later generations, or vice versa. When the discount factor β is exactly $1/\rho$, the planner's only remaining goal is to allocate risk efficiently. Thus, only this discount factor replicates the equilibrium that generations would choose on their own in the original position. In this way, the Rawlsian equilibrium pins down the discount rate in the Ramsey model without invoking intergenerational altruism.

VI. Implications for the Design of Social Security

So far, we have considered how optimal intergenerational risk sharing, as modeled by complete contingent-claims markets, affects the allocation of resources. We now move closer to issues of policy and consider what institutions might support this optimal allocation. The natural institution to consider is social security, because it takes resources from some generations and gives resources to others, which is what is needed to share generational risk. But how should we design a social security system if our

goal is to implement the allocation of resources in the Rawlsian equilibrium?

The first result concerning social security design follows naturally from the results we have already seen:

Proposition 1: Without government intervention, the economy cannot reach the optimal allocation of risk across generations. Similarly, a social security system that relies completely on private accounts also fails to allocate risk optimally.

The first part of the proposition states that the Hobbesian equilibrium is not the same as the Rawlsian equilibrium, which we established in the preceding sections. The second part of the proposition follows for the same reason. In this model, a social security system with private individual retirement accounts does not move the allocation of risk away from the Hobbesian equilibrium: each generation still bears the full risk of shocks to the capital return rather than sharing the risk with other generations. Similar results about the sub-optimality of the equilibrium without intervention are presented by authors such as Bohn (1998) and Rangel and Zeckhauser (1998).

Although it is easy to see that a privatized social security system does not implement the Rawlsian equilibrium, it is less obvious how to describe policies that do. In overlapping-generations models, the government can often achieve the same

allocation of resources in several equivalent ways. For example, a tax or transfer can occur when a person is young or old; with appropriate discounting, this does not matter for the resulting allocation of consumption. For concreteness and realism, we focus on policies that resemble social security systems: the young pay taxes based on their wages, and the old receive transfers. We examine two ways to implement the Rawlsian equilibrium, as expressed in equation (8).

The first is described in this proposition:

Proposition 2: The government can implement the Rawlsian equilibrium using a fully-funded social security system with a trust fund invested in equity claims to capital. The social security benefit responds positively to the capital return, and it follows a random walk.

The proof is straightforward. In essence, the government here takes over the economy and enforces the Rawlsian allocation. It taxes 100 percent of wages (recall that there is no first-period consumption), invests the tax revenue in capital, and then pays out a benefit determined by equation (8). Providing a social security benefit equal to consumption in the Rawlsian equilibrium ensures that the system replicates that equilibrium. This system is feasible because the Rawlsian equilibrium is feasible.

This social security system may seem remote from real-world

policy, but there is another, more natural way to describe this system. The tax rate is constant, the system is fully-funded and invested in equity, and the benefit rises or falls as the economy realizes shocks. In each period, the benefit is based on the system's "permanent income." That is, the benefit is set at a level that could remain constant if there were no more shocks. Seen in this light, the system resembles some proposals for social security reform, which often involve adjusting benefits in response to changes in the system's expected resources.

In the system just described, the social security trust fund must be invested in equity claims to capital. There is, however, another way to reach the Rawlsian allocation that does not require the trust fund to hold equity claims:

Proposition 3: The government can implement the Rawlsian equilibrium using a fully-funded social security system invested in riskless bonds. In this system, the benefits received by the old are negatively indexed to the current return to capital.

To establish this proposition, imagine we were in a world described by Proposition 2, where the trust fund is invested in equity claims to capital. Then suppose the government makes three changes. First, it sells its equity claims to the private sector. Second, it uses the proceeds from that sale to buy riskless debt from the private sector. Third, it adjusts the social security

benefit to insure the private sector against the uncertainty inherent in holding the equity claims. Meanwhile, the private sector engages in the opposite transaction: it buys the equity claims with the proceeds from its debt sale.

In this new system, the social security benefit is consumption in the Rawlsian equilibrium, as determined in equation (8), minus $\theta_t x$. This last term represents the negative indexation: it offsets the gain or loss that the private sector experiences from issuing riskless bonds and buying risky capital. (Note that under our small-shock assumption, riskless bonds pay ρ , while risky capital earns $\rho + \theta_t x$.) This scheme reaches the same allocation of consumption and risk as in Proposition 2--the Rawlsian allocation. But the asset allocation has changed: risky capital is now held in the private sector, rather than by the government.

The message of Propositions 2 and 3 can be summarized as follows. In the Hobbesian equilibrium, capital risk in any period falls entirely on the generation that is old in that period. To move toward the Rawlsian equilibrium, a social security system has to share that risk with future generations. There are two ways to do this. The social security system can hold the economy's capital stock and the risks associated with it. Or the social security system can insure generations for the capital risk they bear through negative indexation.

VII. Two Extensions

This section considers two ways to generalize our model. First, we allow wages as well as capital returns to be uncertain. Second, we assume that agents consume in both periods of their lives. The derivations parallel those for our basic model, so we only sketch the analysis.

A. Wage-Growth Uncertainty

In reality, wages as well as capital returns are uncertain. A generation can suffer bad luck in the form of low wage growth, such as the productivity slowdown from 1974 to 1996. Here we examine intergenerational risk-sharing in the presence of wage as well as asset-return uncertainty.

Let W_t denote the wage in period t and $H_t = W_t/W_{t-1}$ denote the gross growth rate of the wage. H represents the growth in the value of the human capital with which a person is endowed at birth. As a first step toward realism, we allow the mean of H , which we call γ , to be greater than one. Paralleling our assumption about R , we assume that H_t has a two-point distribution: H_t equals $\gamma + z_t$ with probability one half and $\gamma - z_t$ with probability one half, where z_t is a positive constant. Wage-growth H_t is serially uncorrelated, so the level of the wage follows a random walk with drift; this seems realistic as a first approximation. We assume $\gamma + z_t < \rho - x_t$, which assures dynamic efficiency in Diamond's sense. We make no assumption about the contemporaneous correlation of H_t and R_t .

In this version of the model, the consumption of generation t when old is $W_t R_{t+1}$ in the Hobbesian equilibrium. It is the product

of a random-walk variable (the wage) and a serially uncorrelated variable (the return to capital). Because there are two possible realizations of R_t and two of H_t , there are a total of four possible states each period. This means there are 4^t possible histories of the economy through period t . In the original position, there is a market for consumption in each period in each of these histories.

The analysis of this model parallels the previous discussion. One can first solve the model for the case of a wage-growth shock in a single period to see how the risk from this shock is shared among generations. One can then derive a first-order approximation to the general case of wage-growth and capital-return shocks in each period. If uncertainty is constant over time (that is, if $x_t=x$ and $z_t=z$ for all t), then the approximate solution for consumption in the Rawlsian equilibrium is

$$(10) \quad c_t = \rho\gamma^t [1 + (N_t^{\text{RG}} - N_t^{\text{RB}})x(\rho-\gamma)/\rho^2 + (N_t^{\text{HG}} - N_t^{\text{HB}})z/\rho],$$

where N_t^{RG} is the number of periods in which the realization of the R-shock is good and the other N's are defined similarly.

To interpret this expression, note that $\rho\gamma^t$ would be the consumption of generation t when old if wage growth and the capital return were always equal to their means of γ and ρ . Shocks to R and H cause random-walk movements in consumption relative to this baseline. Each shock to the capital return permanently raises or lowers consumption by a fraction $x(\rho-\gamma)/\rho^2$; each shock to wage growth raises or lowers consumption by a fraction $z(1/\rho)$. Consumption follows a random walk, with the innovation in each

period depending on both of the shocks. Generalizing another earlier result, we can show that this consumption behavior is the same as would be chosen by a Ramsey social planner with discount factor γ/ρ .

As in our earlier analysis, moving from the Hobbesian to the Rawlsian equilibrium reduces the impact of shocks on the generations who receive them. In the Hobbesian equilibrium, a positive shock to the capital return at t raises the consumption of the old at t by a fraction x/ρ . In the Rawlsian equilibrium, the effect is $x(\rho-\gamma)/\rho^2$, which is a fraction $(\rho-\gamma)/\rho < 1$ of the effect in the Hobbesian equilibrium. Similarly, a positive wage shock raises the consumption of the generation that receives it by z/γ in the Hobbesian equilibrium, and z/ρ in the Rawlsian equilibrium; the latter is a fraction $\gamma/\rho < 1$ of the former.

In the complete markets of the original position, the two kinds of risk are shared differently. As before, a generation that receives a capital-return shock shares the risk with future generations. However, the original position does not create opportunities to share wage-growth risk with future generations. Because the wage follows a random walk, a wage-growth shock at t already has a proportional effect on all generations born at t and later in the Hobbesian equilibrium, leaving no room for additional risk-sharing. There is, however, an opportunity to share the risk from a wage-growth shock at t with the generation that is old during that period. In the Hobbesian equilibrium, this generation

is unaffected by the shock, because its wage was determined in the previous period. By contrast, in the Rawlsian equilibrium, a wage-growth shock at t affects the consumption of the old at t ; this allows a smaller effect on the young at t and all later generations.

The addition of wage uncertainty has little effect on the nature of the social security systems that implement the Rawlsian equilibrium. Propositions 1-3 of the previous section still hold. The only difference is that both shocks affect the system's resources and hence cause random-walk movements in the level of social security benefits.

B. Consumption in Both Periods of Life

We now relax the assumption that individuals consume only when old. In particular, we assume that an individual born in period t receives utility of

$$(11) \quad U_t = \alpha \log(c_t^y) + (1-\alpha) \log(c_{t+1}^o),$$

where c_t^y and c_t^o are the consumption of the young and the old in period t . Our earlier model is the special case in which $\alpha=0$.

Once again, the Hobbesian equilibrium is simple to derive. The assumption of log utility leads to the result that an agent saves a fraction $1-\alpha$ of his wage. Consumption when young is αW_t , and consumption when old is $(1-\alpha)W_t R_{t+1}$.

To derive the Rawlsian equilibrium, we continue to assume that there are two possible values of R_t and two possible values of H_t for each t . Thus, there are again 4^t history-contingent

consumption goods for each t . In this case, however, an agent born at t consumes goods dated at both t and $t+1$.

At an intuitive level, it is easy to see the Rawlsian equilibrium that arises in this setting. Each agent smooths consumption over the two periods of his life; thus, a shock at t has the same proportional effect on c^y and c^o for all generations born at t and later. Along with this smoothing across periods for each generation, we have the same smoothing across generations as before using the complete contingent-claims markets.

One detail is that a shock to wage growth or the capital return in period t affects both c^y and c^o for generations born at t and later but affects only c^o for the generation born at $t-1$, because c_{t-1}^y is determined before period t . The Appendix gives formulas for c_t^y and c_t^o . The key qualitative features are that the ratio of c_t^y to c_t^o is a constant, and that c_t^y and c_t^o each follow a random walk, rising or falling each period in response to current shocks to wage growth and the capital return. There is perfect risk-sharing both across the old and young alive at the same time and across different periods.

Turning to the implementation of the Rawlsian equilibrium, we find that the spirit of our earlier results continues to hold. The only qualification to Propositions 1-3 concerns the behavior of taxation. Consider Proposition 2, which describes a fully-funded trust fund that holds equity claims to capital. Once again, the social security benefit equals consumption when old in the Rawlsian

equilibrium. Yet, because agents now consume when young, taxes on the young are no longer 100 percent. Instead, taxes are set such that the after-tax wage equals consumption when young in the Rawlsian equilibrium.

In this system, the key parameters are again adjusted in the spirit of the permanent income hypothesis, in the sense that shocks affect the system's resources, leading to permanent changes in benefits and taxes. A good shock to the capital return raises benefits and reduces taxes. A good shock to wage growth raises both benefits and taxes. (This last result ensures that some of the windfall to the young is taxed away to be shared with the current old). The system's parameters are always adjusted in a way that maintains a fixed ratio of the benefit to the after-tax wage. Both the level of taxes and the social security benefit follow a random walk.

VIII. Conclusion

This paper has explored an approach to analyzing intergenerational risk sharing. According to this approach, policymakers designing institutions that share generational risk should attempt to achieve the allocation that the various generations would reach on their own if they could have traded in complete contingent-claims markets. That is, policy should achieve what the invisible hand would if it could.

This approach can be used not only for deriving the optimal

allocation of consumption but also as a guide for the design of a social security system. An obvious but important result from our analysis is the suboptimality of private retirement accounts--a possible social security reform that has received much attention in recent years. Private retirements accounts merely replicate the equilibrium without any intergenerational risk sharing. That is, private retirement accounts leave all generations facing more risk than they should.

Another robust conclusion from our analysis is that the government should spread capital risk among generations in a way that appears absent from current policy. If equity claims to capital are held privately, as they are now, then optimal intergenerational risk sharing requires that social security benefits be negatively indexed to the capital return: social security benefits should be cut when the stock market is doing well. In the absence of such negative indexation, the government should invest the social security trust fund directly in capital. Negative indexation and government ownership of capital seem to be the only mechanisms that allow current capital risk to be shared optimally with future generations.

Several recent proposals for social security reform have, in fact, included such provisions. The Clinton administration, for instance, proposed investing the social security trust fund in equities, as envisioned in our Proposition 2. The negative indexation of benefits to equity returns may seem less likely, but

in fact it is part of the Feldstein proposal for social security reform (see, for example, Feldstein and Samwick, 1999, and Feldstein and Rangelova, 2001). In this plan, individuals would have private accounts invested in capital markets; the more they earn in these accounts, however, the less they would receive in supplemental benefits. This "clawback" provision, as it is often called, resembles the negative indexation envisioned in our Proposition 3. Either approach could implement the Rawlsian equilibrium, raising the expected welfare of all generations. In theory, intergenerational risk sharing offers the prospect of a free lunch.

Admittedly, given economists' limited understanding of these issues, it may be too early to jump to policy conclusions. Even with the extensions in Section VII, the model in this paper makes many strong assumptions: individuals within a generation are homogeneous, wages and capital returns are exogenous, all generations are the same size, and so on. Addressing real-world issues of social security reform will require relaxing these assumptions. Fortunately, the concept of a Rawlsian equilibrium--the equilibrium in an overlapping-generations model with complete contingent-claims markets--is quite general. Future work could investigate the nature of the Rawlsian equilibrium and the institutions that can implement it in a richer variety of settings.

APPENDIX

This Appendix presents details of our analysis that are omitted from the text.

Equilibrium Prices with a Single Capital-Return Shock

Here we derive the Rawlsian equilibrium when there is a single capital-return shock in period j . We do this by deriving two necessary conditions for the equilibrium, equations (2) and (3).

Equation (2) states that, for $t \geq j$, the relative price of consumption at t and $t+1$ in the same history must equal the gross interest rate ρ . To see why this condition must hold, suppose first that $P_{ti}/P_{(t+1)i} < \rho$. In this case, an agent born at t can buy a unit of good ti , save it to acquire ρ units of good $(t+1)i$, sell enough of good $(t+1)i$ to pay for his purchase of good ti , and still have some left over. This arbitrage possibility would create an infinitely large demand for good ti , which could not be satisfied by the finite supply.

Now suppose that $P_{ti}/P_{(t+1)i} > \rho$. In this case, no agent born at t will save, because he can obtain a higher return by selling his endowment of good ti and buying good $(t+1)i$. Thus the capital stock is zero in period t . This cannot be an equilibrium, because no capital in period t means there are not enough resources from $t+1$ onward to make every generation as well off as in the Hobbesian equilibrium. If some generation is worse off than in the Hobbesian equilibrium, the allocation cannot be the Rawlsian equilibrium, because agents have the option of avoiding the markets and

receiving their Hobbesian consumption levels.¹

Now consider equation (3), which gives the relative price of consumption in the good and bad histories. This equation follows from two underlying conditions. The first is a first order condition for utility maximization: the relative price Q must equal the ratio of marginal utilities of consumption in the good and bad histories. With log utility, the ratio of marginal utilities is the inverse of the ratio of consumption levels. Thus $Q = c_{tG}/c_{tB}$ for $t \geq h$.

The other condition underlying (3) is that, in each possible history, the present value of total consumption beginning in period j must equal the present value of resources beginning at j , given the gross interest rate ρ that holds from $j+1$ on. The present value of resources is the gross return on capital at j plus the present value of wages at $j, j+1, \dots$. If the present value of consumption were greater than the present value of resources, the allocation would not be feasible. If the present value of resources were larger than the present value of consumption, the allocation would be inefficient, and hence could not be a Walrasian equilibrium.

In the good history, the gross capital return in period j is

¹ If the capital stock is zero in period t , the economy's resources from $t+1$ on are given by the certain wage of one at $t+1, t+2, \dots$. In the Hobbesian equilibrium, consumption is $\rho > 1$ at $t+1, t+2, \dots$. Thus there are not enough resources to give each generation its consumption level in the Hobbesian equilibrium with certainty. And one can show that randomization would only make matters worse, because it would create uncertainty about consumption without raising its average level.

$\rho + x_j$ and the wages at $j, j+1, \dots$ are $1, 1, \dots$. The present value of these resources is $\rho + x_j + \rho/(\rho-1)$. In the bad history, the wages are the same but the capital return at j is $\rho - x_j$; the present value of resources is $\rho - x_j + \rho/(\rho-1)$. The ratio of these two present values must equal the ratio of the present values of consumption in the two histories. Recall that the ratio of consumption in the two histories is Q in each period; this implies that the ratio of present values of consumption is Q . Setting Q equal to the ratio of present values of resources yields equation (3).

Equilibrium Consumption Levels with a Single Capital-Return Shock

Given the relative prices in equations (2)-(3), one can derive equilibrium consumption levels from agents' utility-maximization problems. A generation born in period $t \geq j$ has an endowment of one unit of good tG and one unit of good tB , because his wage is one in both histories. If we treat good tG as the numeraire, the value of an agent's endowment at t is $(1+Q)$. The agent wishes to consume in period $t+1$ in each history. Given the relative prices in (2) and (3), the agent's budget constraint is

$$(A1) \quad c_{(t+1)G} + Qc_{(t+1)B} = \rho(1+Q).$$

Maximizing the average of utility, $\log(C_{t+1})$, over the two histories subject to (A1) yields the solutions for $c_{(t+1)G}$ and $c_{(t+1)B}$ in equation (5).

The generation born at $j-1$ and old in period j receives a unit

of good $j-1$ (this is not indexed by G or B because the shock has not yet occurred). By saving, he can transform his unit of good $j-1$ into $\rho+x_j$ units of good jG and $\rho-x_j$ units of good jB . He wishes to consume goods jG and jB , and faces the budget constraint

$$(A2) \quad c_{jG} + Qc_{jB} = \rho + x_j + Q(\rho-x_j) .$$

Maximizing expected utility subject to (A2) yields the consumption levels in (4).

The Case of Many Small Shocks

Equation (6) gives a first order approximation in x_j of equations (4) and (5), the equilibrium consumption levels in the example of a single capital-return shock. To see how (6) is derived, consider c_{tG} for $t \geq j+1$. Evaluating the expression in (5) at $x_j=0$ yields $c_{tG}=\rho$, since $Q=1$ when $x_j=0$. Differentiating with respect to x_j yields

$$(A3) \quad dc_{tG}/dx_j = (\rho/2)(dQ/dx_j) .$$

Taking the derivative of Q (equation (3)) and evaluating it at $x_j=0$ yields

$$(A4) \quad dQ/dx_j|_{x_j=0} = 2(\rho-1)/\rho^2 .$$

Substituting (A4) into (A3) yields

$$(A5) \quad dc_{tG}/dx_j|_{x_j=0} = (\rho-1)/\rho .$$

This result leads to the approximate solution for c_{tG} in (6). The results for c_{tB} and for consumption at $t=j$ are obtained similarly.

Our use of first-order approximations makes it easy to go from one shock to the general case of $x_t > 0$ for all t . In the general case, we are interested in deriving

$$c_{th}(x_1, \dots, x_t)$$

for all periods t and all histories h through t . A first-order approximation in x_1, \dots, x_t yields

$$(A6) \quad c_{th} \approx c_{th}(0, \dots, 0) + \sum_{s=1}^t \partial c_{th} / \partial x_s(0, \dots, 0) x_s .$$

In this expression, the first term on the right is ρ , the equilibrium consumption level when there are no shocks. Within the sum, a term $\partial c_{th} / \partial x_s(0, \dots, 0)$ can be determined as follows. Consider $c_{th}(0, \dots, 0, x_s, 0, \dots, 0)$, i.e. c_{th} as a function of x_s when all other x 's are zero. Differentiating this function with respect to x_s and evaluating it at $x_s=0$ yields $\partial c_{th} / \partial x_s(0, \dots, 0)$. But $c_{th}(0, \dots, 0, x_s, 0, \dots, 0)$ is just the solution for conditional consumption in period t when there is a single capital-return shock at s . This solution is given in equations (4) and (5). In deriving (6), we used the results that the derivatives of (4) and (5) evaluated at $x_s=0$ are $(\rho-1)/\rho$ when the capital-return shock at s is good and $-(\rho-1)/\rho$ when the shock is bad. These results imply $\partial c_{th} / \partial x_s(0, \dots, 0) = \theta_s(\rho-1)/\rho$. Substituting this result into (A6) yields equation (7), the approximate solution for consumption in the general case.

The Rawlsian Equilibrium in the General Case

Here we provide approximate solutions for consumption in the Rawlsian equilibrium when agents consume in both periods of life and both wage growth and the capital return are uncertain. We assume constant levels of uncertainty: $x_t=x$ and $z_t=z$ in all periods.

The results are generalizations of equations (8) and (10) and are derived in the same way as those equations:

$$(A7) \quad c_t^i = c_t^{i*} [1 + (N_t^{RG} - N_t^{RB})Kx + (N_t^{HG} - N_t^{HB})K'z],$$

$$K = (1-\alpha)(\rho-\gamma) / [\rho^2 - \alpha\rho(\rho-\gamma)], \quad K' = 1 / (\rho - \alpha\rho + \alpha\gamma),$$

where c_t^i for $i=y,o$ are consumption of the young and the old and c_t^{i*} is consumption in the certainty case of $R_t=\rho$ and $H_t=\gamma$ for all t ($c_t^{y*}=\alpha\gamma^t$ and $c_t^{o*}=(1-\alpha)\gamma^t\rho$). As discussed in the text, the ratio of c_t^y to c_t^o is a constant equal to $c_t^{y*}/c_t^{o*} = \alpha/[(1-\alpha)\rho]$. Both c_t^y and c_t^o follow random walks in response to shocks to H and R .

References

- Abel, Andrew B., N. Gregory Mankiw, Lawrence H. Summers, and Richard J. Zeckhauser, "Assessing Dynamic Efficiency: Theory and Evidence," Review of Economic Studies 56 (January 1989), 1-20.
- Barro, Robert J., "Are Government Bonds Net Wealth," Journal of Political Economy 82 (November/December 1974), 1095-1117.
- Bohn, Henning, "Risk Sharing in a Stochastic Overlapping Generations Economy," University of California at Santa Barbara, December 1998.
- Diamond, Peter A., "National Debt in a Neoclassical Growth Model," American Economic Review 55 (December 1965), 1126-1150.
- Feldstein, Martin and Elena Rangelova, "Individual Risk in an Investment-Based Social Security System," NBER Working Paper #8074, 2001 (American Economic Review, forthcoming).
- Feldstein, Martin, and Andrew Samwick, "Social Security Benefits Through Personal Retirement Accounts: An Update Based on the 1998 Social Security Trustees Report," NBER Working Paper #6540, March 1999.
- Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy 86 (December 1978), 971-987.
- Rangel, Antonio, and Richard Zeckhauser, "Intergenerational Risk Sharing Properties of Market and Political Institutions," Stanford University, November 1998
- Rawls, John, A Theory of Justice, Harvard University Press, 1971.
- Shiller, Robert, "Social Security and Intergenerational, Intragenerational, and International Risk Sharing," Carnegie-Rochester Conference Series on Public Policy 50 (June 1999), 165-204.
- Smetters, Kent, "The Taxation of Risky Capital Income in an Overlapping Generations Stochastic Production Economy," Wharton School, March 2000.

