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

Citation

Ball, Laurence and N. Gregory Mankiw. 2007. Intergenerational risk sharing in the spirit of Arrow, Debreu, and Rawls, with applications to social security design. *Journal of Political Economy* 115(4): 523-547.

Published Version

<http://dx.doi.org/10.1086/520646>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3443106>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

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

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 is a revision of NBER Working Paper 8270. We are grateful for suggestions from Christopher Carroll, David Romer, Stephen Zeldes, the referee, and the editor.

[*Journal of Political Economy*, 2007, vol. 115, no. 4]
© 2007 by The University of Chicago. All rights reserved. 0022-3808/2007/11504-0001\$10.00

This paper explores a deviation from Arrow-Debreu theory that arises from the 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 (2001), and Smetters (2002). 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?¹

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 in which 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 policy maker might try to use these instruments to achieve the allocation of risk that the invisible hand would reach if it could. That is, the policy maker 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

¹ Bohn's paper is perhaps the closest to this one. Bohn's model is similar in spirit to ours, and he shows how various fiscal policies shift risk between generations. However, Bohn does not try to determine the optimal allocation of risk or the social security systems that implement this allocation.

social security trust fund should be invested and how 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 *natural equilibrium* because it is the equilibrium that arises under realistic assumptions about trading possibilities.

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 *complete-markets equilibrium*.

Sections IV and V solve for the complete-markets 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 many generations face 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 natural equilibrium, where consumption is serially uncorrelated from generation to generation, consumption in the complete-markets equilibrium follows a random walk.

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 complete-markets 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 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. Policy makers can also implement the complete-markets 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 considers some variations on our basic model, and Section VIII presents conclusions.

II. The Model and the Natural Equilibrium

This section describes the basic version of our overlapping-generations model, which we keep simple to build intuition. A fixed number of people are born each period, and everyone lives for two periods. When young, a person receives an endowment of one unit of the consumption good. Each person consumes only when he is old. Thus a person saves his entire endowment, and it becomes next period's capital stock. He consumes the return on this saving, including the principal, when he is old. There are no bequests. We assume log utility over consumption; thus expected lifetime utility for an individual born in period t is

$$E[U_t] = E[\ln(c_{t+1})], \quad (1)$$

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.

The economy contains many competitive firms. Firms have a linear production function

$$Y_t = R_t K_t, \quad (2)$$

where Y is output and K is capital. Firms obtain capital from the old. Capital depreciates 100 percent at the end of the period. These assumptions imply that the gross return on savings at period $t - 1$ is R_t . The firms make zero profits, so it does not matter who owns them.

We take the capital return R_t to be an exogenous random variable. It is distributed independently over time and has a two-point distribution. Let μ be the average value of the interest rate. The return R_t equals $\mu + \sigma_t$ with probability one-half and $\mu - \sigma_t$ with probability one-half. The term σ_t measures the degree of dispersion in the capital return in period t . It is natural to focus on the case in which σ_t is time-invariant (and we will do so below), but it will prove useful to have in hand the general case in which σ_t varies across periods. We assume that the lowest possible return, $\mu - \sigma_t$, is greater than one; this ensures 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 *natural equilibrium*.

Given the environment just described, the natural equilibrium is trivial to derive. The generation born at t saves its endowment and consumes its wealth when old. Thus its consumption, c_{t+1} , equals R_{t+1} . Given the two possible realizations of R_{t+1} , the generation's expected utility is $\frac{1}{2} [\ln(\mu + \sigma_{t+1}) + \ln(\mu - \sigma_{t+1})]$.

Because R_{t+1} is uncorrelated over time, each generation bears all of a single idiosyncratic risk—the return risk in the period in which it happens to be old. Note that consumption is independently distributed from generation to generation.

This natural equilibrium is clearly inefficient. There would be gains if the old generation 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 Complete-Markets 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 0, 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. We let θ_t denote the state of the economy in period t . The state captures the period's capital return: it is “high” ($\theta_t = H$) if $R_t = \mu + \sigma_t$ and “low” ($\theta_t = L$) if $R_t = \mu - \sigma_t$. We let Θ denote the set of possible states: $\Theta = \{H, L\}$. A “history” of the economy through period t is a sequence of states for periods 1, ..., t . We denote a history by $\theta^t = (\theta_1, \dots, \theta_t)$. There are 2^t possible histories of the economy through period t , each of which occurs with probability $(\frac{1}{2})^t$.

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 through that period. For each period t , there are 2^t markets for contingent consumption goods indexed by θ^t .

A person born in period t receives an endowment of one unit of the consumption good in all histories of the economy through t . Thus his endowment is one unit of good θ^t for all possible θ^t . The consumer sells this endowment and uses the proceeds to buy conditional consumption goods at $t + 1$, when he wants to consume. He maximizes his expected utility over all possible histories subject to his budget constraint. With $p(\theta^t)$ the price of the consumption good in period t and history θ^t , the individual's problem is

$$\begin{aligned} \max_{\theta^{t+1} \in \Theta \times \dots \times \Theta} & \sum u(c_{t+1}(\theta^{t+1}))[(\frac{1}{2})^t] \\ \text{subject to} & \sum_{\theta^{t+1} \in \Theta \times \dots \times \Theta} c_{t+1}(\theta^{t+1})p(\theta^{t+1}) = \sum_{\theta^t \in \Theta \times \dots \times \Theta} p(\theta^t), \end{aligned}$$

where $u(c_{t+1}) = \ln(c_{t+1})$.

Firms also participate in the contingent-claims markets. A firm buys contingent goods at t , uses them as capital at $t + 1$, and sells its output. If a firm purchases good θ^t , it produces two contingent goods dated at $t + 1$ —one for each of the histories that is a continuation of θ^t . We denote these two histories by (θ^t, θ_{t+1}) for $\theta_{t+1} = H, L$. With a unit of good θ^t , the firm produces $\mu + \sigma_{t+1}$ units of good (θ^t, H) and $\mu - \sigma_{t+1}$ units of good (θ^t, L) . It chooses its purchases of goods, $i(\theta^t)$, to maximize profits:

$$\begin{aligned} \max_{i(\theta^t)} & -i(\theta^t)p(\theta^t) + i(\theta^t)[(\mu + \sigma_{t+1})p(\theta^t, H) + (\mu - \sigma_{t+1})p(\theta^t, L)] \\ & \forall \theta^t \in \Theta \times \dots \times \Theta. \end{aligned}$$

The maximand will turn out to be zero in equilibrium.

A Walrasian auctioneer finds the set of equilibrium prices $p(\theta^t)$ that equates supply and demand for each history-contingent good. Demand for a good is the sum of consumption and investment, and supply is the endowment of the young plus the return on the previous period's investment. Thus the market-clearing conditions are

$$i_t(\theta^t) + c_t(\theta^t) = 1 + i_{t-1}(\theta^t)R_t(\theta^t)$$

for all periods and histories.

It is easy to derive one condition for equilibrium prices. The firm's objective function is linear in $i(\theta^t)$, its purchase of good θ^t . For the firm to buy a positive but finite amount of the good (which it will in equi-

librium), the derivative of profits with respect to $i(\theta')$ must be zero. This requires that prices satisfy

$$p(\theta') = (\mu + \sigma_{t+1})p(\theta', H) + (\mu - \sigma_{t+1})p(\theta', L). \quad (3)$$

The left side of this expression is the cost of buying a unit of good θ' and the right is the firm's revenue from buying the unit.

Beyond this point, finding the complete-markets equilibrium 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 consider a case with many shocks but solve the model using a first-order approximation. Thus we derive a solution that is valid when shocks to capital returns are small. Each of these two special cases yields its own insights into the nature of the complete-markets 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, $\sigma_j > 0$ and $\sigma_t = 0$ 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 Complete-Markets Equilibrium

In this example, the possible histories of the economy collapse to a simple set. In period j , there are two possible states: the state in which $R_t = \mu + \sigma_t$ (H) and the state in which $R_t = \mu - \sigma_t$ (L). In all other periods, there is only a single state with $R_t = \mu$. There is one possible history of the economy through t for $t < j$ and two possible histories through $t \geq j$: the history with state H at j and the history with state L at j . Figure 1 shows the possible histories in a graph. We index the consumption goods in these histories by t for $t < j$ and by tH , tL for $t \geq j$.

The Appendix describes in detail the solution for the complete-markets equilibrium. Here, we sketch the main steps in the analysis.

For periods before j , when the shock occurs, consumption in the complete-markets equilibrium is the same as in the natural equilibrium. Consumption before j cannot be made contingent on the shock. To find consumption for $t \geq j$, we first find the prices of contingent goods

FIG. 1.—History-contingent consumption (case of one shock)

in those periods, which we denote by P_{iH} and P_{iL} . The Appendix shows that these prices satisfy

$$\frac{P_{ii}}{P_{(t+1)i}} = \mu \quad \text{for } i = H, L \quad (4)$$

and

$$\frac{P_{iL}}{P_{iH}} = Q \equiv \frac{\mu^2 + (\mu - 1)\sigma_j}{\mu^2 - (\mu - 1)\sigma_j}. \quad (5)$$

These equations define all relative prices for goods dated j or later.

These equilibrium conditions are simple to interpret. Condition (4) concerns the prices of the good in different periods but the same realization of history. It is a special case of the no-arbitrage condition for firms, equation (3). It says that the relative price of consumption at t and $t + 1$ must equal the return from investing at t . In this example, the return is fixed at μ for $t \geq j$.

Condition (5) gives the relative price of consumption in the high- and low-return histories, which is the same for all $t \geq j$. The key result is that $Q > 1$, which follows from our assumptions that $\mu > 1$ and $\sigma_j > 0$. This result means that it costs more than one unit of consumption in the high-return history to buy a unit in the low-return history. This feature of relative prices is necessary to induce agents to consume more when more resources are available.

Given the equilibrium relative prices, the Appendix derives the history-contingent consumption of each generation. For the generation born at $j - 1$ (the one that experiences the shock), consumption when old is given by

$$\begin{aligned} C_{jH} &= \frac{(1 + Q)\mu + (1 - Q)\sigma_j}{2}, \\ C_{jL} &= \frac{(1 + Q)\mu + (1 - Q)\sigma_j}{2Q}. \end{aligned} \quad (6)$$

For all generations born at j and later, consumption is

$$\begin{aligned} C_{iH} &= \frac{\mu(1 + Q)}{2}, \\ C_{iL} &= \frac{\mu(1 + Q)}{2Q}, \quad t \geq j + 1. \end{aligned} \quad (7)$$

Equations (4)–(7) fully describe the complete-markets equilibrium.

B. Discussion of the Complete-Markets Equilibrium

The solution we have just described has two notable properties. First, the ratio of consumption in the high- and low-return histories equals $Q > 1$ 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 natural 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 $(\mu + \sigma_i)/(\mu - \sigma_i)$, which is greater than Q . Thus this generation reduces its risk by moving from the natural equilibrium to the complete-markets equilibrium, where it can share risk with future generations.

The second notable result concerns *average* consumption in the natural and complete-markets equilibria. For generations born at j and later, average consumption over the high- and low-return histories is $\mu[(1 + Q)^2]/[4Q]$. This exceeds μ , which is these generations' consumption in the natural equilibrium. Thus average consumption is higher in the complete-markets equilibrium than in the natural equilibrium for all these generations. A bit more algebra shows that for the generation born at $j - 1$, average consumption is lower in the complete-markets equilibrium. Of course, for all generations, expected utility must be higher in the complete-markets equilibrium, for the natural equilibrium is still budget feasible.

These results have a simple interpretation. In the natural 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 high-return history than they give up in the low-return history.

V. The Case of Many Shocks

Having explored the special case of a single shock, we now examine a more general case in which there are shocks in many periods. For some period T , we assume $\sigma_t > 0$ for all $t \leq T$ and $\sigma_t = 0$ for $t > T$. The value of T can be any positive integer, so our analysis covers any finite number of shocks, however large. In the original position, there are now 2^t markets for history-contingent consumption in each period $t \leq T$ and 2^T markets for $t > T$.

In this case we cannot find an exact solution for the complete-markets equilibrium, but we can learn about it in two ways. First, we present qualitative results about the equilibrium and how it differs from the natural equilibrium. Second, we use a first-order approximation to solve for consumption in the complete-markets equilibrium. This approximation is valid when shocks to capital returns are small.

A. General Results

In our one-shock example, a key property of the complete-markets equilibrium is that the effect of a shock is spread across all current and future generations. This property carries over to the more general model. The following proposition states this formally.

PROPOSITION 1. Consider any periods $\tau < T$ and $t \geq \tau$. Let θ' be any history through t in which $\theta_\tau = H$, and let θ'' be the history that is the same as θ' except that $\theta_\tau = L$. Then $c_t(\theta') > c_t(\theta'')$.

In other words, with the shocks in other periods held constant, a high rather than low return in period τ raises consumption at τ and all later periods. The risk from a shock in any period is spread into the future. (Proposition 1 is proved in the Appendix.)

In the one-shock example, moving from the natural equilibrium to the complete-markets equilibrium makes all generations better off. This result also generalizes to the case of many shocks, as in the following proposition.

PROPOSITION 2. For every generation, expected utility in the complete-markets equilibrium is higher than expected utility in the natural equilibrium.

In other words, opening up the markets in the original position leads to a Pareto improvement.

This proposition follows from proposition 1. That result implies that each generation's consumption in the complete-markets equilibrium differs from its consumption in the natural equilibrium, in which a shock affects only one generation. In the complete-markets case, it is feasible for each generation to pick the natural allocation: it can eschew the contingent-claims markets, save its endowment, and consume the return on this saving. Since it chooses a different allocation, the one it chooses must provide higher expected utility.

B. An Approximate Solution for Consumption

We would like to find an explicit solution for consumption in the complete-markets equilibrium. To simplify this problem, we use a first-order approximation that is valid as long as the shocks, σ_t , are small.

That is, we derive the complete-markets equilibrium when there are small fluctuations in capital returns.

To understand the approximation, let us write consumption in the complete-markets equilibrium as $c_t(\theta^t; \sigma_1, \dots, \sigma_T)$. Consumption depends on the history θ^t , which tells whether capital returns are “high” or “low” in each period through t . Consumption also depends on the parameters $\sigma_1, \dots, \sigma_T$, which determine the high and low returns. (In period τ , the high return is $\mu + \sigma_\tau$ and the low return is $\mu - \sigma_\tau$.) For a given history θ^t , we take a first-order approximation of $c(\cdot)$ in $(\sigma_1, \dots, \sigma_T)$. We approximate around $(\sigma_1, \dots, \sigma_T) = (0, \dots, 0)$, the case of constant 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 single shock in period $j \leq T$. Equations (6) and (7) give consumption for period j and after in terms of the relative price Q . Substituting the expression for Q , equation (3), into (6) and (7) yields consumption in terms of σ_j and μ . Taking a first-order approximation in σ_j around $\sigma_j = 0$ yields

$$\begin{aligned} C_{tH} &\approx \mu + \left(\frac{\mu - 1}{\mu}\right) \sigma_j, \\ C_{tL} &\approx \mu - \left(\frac{\mu - 1}{\mu}\right) \sigma_j, \quad t \geq j. \end{aligned} \quad (8)$$

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

According to equation (8), the shock to the capital return causes consumption to rise or fall by a fraction $(\mu - 1)/\mu$ of the shock for each generation born at $j - 1$ and later. Note that 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 high- and low-return 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 (8) 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 of the economy, we let D_t be an indicator variable equal to one if the history includes a high capital return in period t ($\theta_t = H$) and minus

one if it includes a low return at t ($\theta_t = L$). In any history, consumption in period t is given by

$$c_t \simeq \mu + \sum_{j=1}^t D_j \left(\frac{\mu - 1}{\mu} \right) \sigma_j. \quad (9)$$

If all the shocks through period T are the same size ($\sigma_j = \sigma$ for all $j \leq T$), then this expression reduces to

$$c_t \simeq \mu + (N_t^H - N_t^L) \left(\frac{\mu - 1}{\mu} \right) \sigma, \quad (10)$$

where N_t^H is the number of periods through t with high capital returns ($R = \mu + \sigma$) and N_t^L is the number with low returns ($R = \mu - \sigma$). A generation's consumption is raised by a fixed amount for every high return in the past and reduced by the same amount for every low return.²

Note that the last equation implies

$$c_t - c_{t-1} \simeq D_t \left(\frac{\mu - 1}{\mu} \right) \sigma, \quad t \leq T. \quad (11)$$

In each period through T , the change in consumption is proportional to the current shock. Thus, even though consumption in the natural equilibrium was serially uncorrelated, consumption in the complete-markets 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.

C. Risk Sharing and Investment

Random-walk consumption is a familiar result. Hall (1978) finds that optimal consumption for an infinitely lived consumer is approximately a random walk. Hall's result also applies to a social planner choosing consumption for different generations. Consumption is approximately a random walk if the planner maximizes the discounted sum of all generations' expected utilities.

In our model, however, random-walk consumption is *not* the choice of any social planner. Instead, it is the equilibrium outcome in markets for risk sharing. The invisible hand of the market spreads risk across

² Another way to write our result is

$$c_t = \mu + \sum D_j \left(\frac{\mu - 1}{\mu} \right) \sigma_j + o(\|\sigma_1, \dots, \sigma_T\|).$$

That is, consumption equals the expression in eq. (9) plus terms that are second-order or higher in the sizes of the shocks.

generations, making consumption a random walk. No policy maker needs to intervene—as long as agents can trade in an original position before the beginning of time.

It is instructive to look behind the contingent-claims markets to see the mechanics of how risk sharing occurs. In a static setting, risk sharing means that consumption is shifted from one agent to another at a point in time. Here, risk sharing occurs between agents alive in different periods and therefore requires shifts in the paths of *aggregate* consumption and investment.

Suppose, for example, that there is a low realization of the capital return in period j . The old at j receive a transfer that partially offsets this shock. This transfer must come from the endowment of the young, so less is left for investment. Lower investment means fewer resources for the young when they become old at $j + 1$, so they, too, receive a transfer from the next generation. In this way, a low realization of the capital return reduces investment throughout the future. A high realization does the reverse. In the natural equilibrium, by contrast, investment is fixed at one, the endowment of the young.

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 complete-markets equilibrium?

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

PROPOSITION 3. Without government intervention, the economy cannot reach the optimal allocation of risk across generations.

This proposition follows from the fact that, in the absence of government intervention, the economy reaches the natural equilibrium. We have seen that this allocation is Pareto inferior to the complete-markets equilibrium. Similar results about the suboptimality of the equilibrium without intervention are presented by authors such as Bohn (1998) and Rangel and Zeckhauser (2001).

Proposition 3 is relevant for evaluating one proposal for social security reform, a system that mandates savings through individual retirement accounts. In our model, such a system does not move the allocation of risk away from the natural equilibrium: each generation still bears the

TABLE 1
A SOCIAL SECURITY SYSTEM INVESTED IN EQUITY

Transactions of the generation born in period t :
 When the generation is young (period t):
 • It pays a tax of one (its entire endowment)
 When the generation is old (period $t + 1$):
 • It receives a social security benefit of

$$c_{t+1}^{CM} \simeq \mu + \sum_{j=1}^{t+1} D_j[(\mu - 1)/\mu]\sigma_j$$

full risk of shocks to the capital return rather than sharing the risk with other generations. Therefore, a social security system based entirely on individual accounts is inefficient.

Although it is easy to see that a privatized social security system does not implement the complete-markets 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, and the old receive transfers. We examine two ways to implement the complete-markets equilibrium. The first is described in the following proposition.

PROPOSITION 4. The government can implement the complete-markets 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 approximately follows a random walk.

The proof is straightforward. In essence, the government here centrally plans the economy. It holds all assets and enforces the complete-markets allocation of consumption.

Table 1 describes how the government treats the generation born in period t . It taxes 100 percent of the generation's endowment, which is one unit of the consumption good (recall that there is no first-period consumption). When the generation is old, at $t + 1$, it receives a social security benefit equal to its consumption in the complete-markets equilibrium (c_{t+1}^{CM}). This rule for benefits ensures that the system replicates the equilibrium. Proposition 1 tells us that the benefit responds positively to the capital return. The benefit is approximated by equation (9), which is a random walk.

In this system, the government invests its tax revenue in capital. In each period, its investment equals i_t^{CM} , the level of investment in the

TABLE 2
A SOCIAL SECURITY SYSTEM INVESTED IN SAFE DEBT

Transactions of the generation born in period t :

When the generation is young (period t):

- It pays a tax of one (its entire endowment)
- It sells a quantity i_t^{CM} of safe debt to the government
- It buys a quantity i_t^{CM} of capital

When the generation is old (period $t + 1$):

- It pays $i_t^{CM}R_{t+1}^*$ to retire its debt
- It earns $i_t^{CM}R_{t+1}$ from its capital
- It receives a social security benefit of

$$c_{t+1}^{CM} - i_t^{CM}(R_{t+1} - R_{t+1}^*) \approx \mu + \sum_{j=1}^{t+1} D_j[(\mu - 1)/\mu]\sigma_j - D_{t+1}\sigma_{t+1}$$

$$\approx \mu + \sum_{j=1}^t D_j[(\mu - 1)/\mu]\sigma_j - (1/\mu)D_{t+1}\sigma_{t+1}.$$

complete-markets equilibrium. The paths of consumption and investment are feasible because the complete-markets equilibrium is feasible.

This social security system may seem remote from real-world policy, but there is another, more plausible, way to describe it. Under 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 as described by equation (9) is based on the system's "permanent income." That is, the benefit is set at a level that could remain constant without further shocks. Seen in this light, the system resembles some proposals for social security reform, which often involve both government ownership of equity and adjustment of benefits 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 complete-markets allocation that does not require the trust fund to hold equity claims.

PROPOSITION 5. The government can implement the complete-markets 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, we construct a social security system with the properties it states. This system is summarized in table 2. As in the previous system, the government taxes 100 percent of the young's endowment. It also buys safe debt from the young. In period t , the government purchases a quantity i_t^{CM} of debt with a gross interest rate of R_{t+1}^* , the equilibrium rate for safe debt maturing at $t + 1$. The young use the proceeds from selling debt to buy a quantity i_t^{CM} of capital.

When a generation is old, it receives the return on its capital, $i_t^{CM}R_{t+1}$. It also pays $i_t^{CM}R_{t+1}^*$ to the government to retire its debt. Finally,

the old receive a social security benefit that equals consumption in the complete-markets equilibrium, c_{t+1}^{CM} , minus a term $i_t^{CM}(R_{t+1} - R_{t+1}^*)$. This term offsets the generation's net gain from being long in capital and short in debt. In effect, the government insures the private sector against the uncertainty it has taken on by holding equity. With this adjustment, the social security benefit and the old's net asset income add up to c_{t+1}^{CM} , implying that the system replicates the complete-markets equilibrium.

Table 2 gives an approximate solution for the social security benefit. The term $R_{t+1} - R_{t+1}^*$ is the difference between the return on equity, $\mu + D_{t+1}\sigma_{t+1}$, and the return on safe debt. One can show that, under the assumption of small shocks, the return on debt equals μ , the average return on equity. (As we have seen, the compensation required to take on risk is second-order.) Therefore, $R_{t+1} - R_{t+1}^*$ is approximately $D_{t+1}\sigma_{t+1}$, which implies that $i_t^{CM}(R_{t+1} - R_{t+1}^*)$ is approximately $D_{t+1}\sigma_{t+1}$. This amount is subtracted from the expression for c_{t+1}^{CM} to give the social security benefit.³

In the solution for the social security benefit, the coefficient on the current shock to the capital return, $D_{t+1}\sigma_{t+1}$, is $[(\mu - 1)/\mu] - 1 = -1/\mu < 0$. Thus, as stated in proposition 5, the benefit is negatively indexed to the current capital return. This indexation is how the government insures the private sector against capital-return risk.⁴

The message of propositions 4 and 5 can be summarized as follows. In the natural equilibrium, capital risk in any period falls entirely on the generation that is old in that period. To move toward the complete-markets 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.

³ Note that $i_t^{CM} = 1 + O(\|\sigma_1, \dots, \sigma_T\|)$; i.e., shocks to capital returns cause first-order changes in investment. However,

$$\begin{aligned} i_t^{CM}(R_{t+1} - R_{t+1}^*) &= [1 + O(\|\sigma_1, \dots, \sigma_T\|)][D_{t+1}\sigma_{t+1} + o(\|\sigma_1, \dots, \sigma_T\|)] \\ &= D_{t+1}\sigma_{t+1} + o(\|\sigma_1, \dots, \sigma_T\|), \end{aligned}$$

justifying the approximation in table 2.

⁴ We have examined the case of small shocks, but the negative-indexation result is general: without any approximation, one can show that the social security benefit responds negatively to the current capital return. This result follows from proposition 1, which states that a higher capital return raises consumption in all future periods. To make it feasible to raise future consumption, the government must increase its current level of assets. When the private sector owns the capital stock, the capital return does not directly affect the government's resources; thus an increase in the government's assets requires a decrease in the current social security benefit.

VII. Extensions

Here we discuss some variations on our model that provide additional insight into the nature of the complete-markets equilibrium and optimal policy.

A. *Two Generalizations*

Two generalizations of our model are straightforward. We summarize them here and provide details in our working paper (Ball and Mankiw 2001).

In the first generalization, agents' initial endowments as well as capital returns are uncertain. This assumption captures uncertainty about workers' wages when young. We assume that the endowment follows a random walk with drift. In the natural equilibrium, a shock to the endowment affects the consumption of the current young and all future generations, because the shock is permanent, but not the consumption of the current old. In the complete-markets equilibrium, the old share the risk from endowment shocks. The optimal social security system is the same as before, except that shocks to endowments as well as capital returns cause changes in the level of benefits.

The other generalization modifies the utility function so that individuals consume in both periods of life. In this case, the complete-markets equilibrium produces perfect risk sharing across agents consuming in the same period: the ratio of consumption by the old and the young is a constant. For the case of small shocks, aggregate consumption follows a random walk, with the consumption of young and old responding together to endowment and capital-return shocks. Once again, the optimal social security system is not greatly affected by the generalization. The only new feature is that the tax rate on endowments is less than one, so the young can consume. Both this tax rate and the level of social security benefits follow random walks.

B. *A Model without Investment*

Risk sharing in our baseline model involves reallocations of consumption over time, which are accomplished through shifts in investment. One might wonder whether such reallocations are essential for the model's contingent-claims markets to raise welfare. To explore this issue, we consider a variation on the model in which it is not technologically feasible to shift aggregate consumption. We find that some but not all of our results are robust. (We only sketch our derivations since they parallel our analysis of the main model.)

Specifically, we consider a model (suggested by the editor) in which

there is no investment. As in our main model, agents have an endowment of one when young. They consume in both periods of life; the period utility function is given by equation (1), and agents maximize a weighted sum of utility in the two periods. All of a period's endowment must be consumed by someone in that period, which implies no investment. However, there is a type of capital: an infinitely lived tree that produces a random amount of the consumption good in each period (as in Lucas [1978]). The tree's output in period t , R_t , equals $\mu + \sigma$ and $\mu - \sigma$ with equal probability. The tree is endowed to the first old generation.

In the natural equilibrium of this model, the old receive the output of the tree and then, having no further use for it, sell it to the young. Thus the ownership of the tree is passed down to each generation. One can show that the equilibrium price of the tree is a constant q that lies between zero and one. In period t , consumption of the young and old are given by

$$c_t^y = 1 - q,$$

$$c_t^o = R_t + q.$$

The young consume the remainder of their endowment after buying the tree, and the old consume the tree's output plus the revenue from selling it.

Note that the consumption of the young is constant and the consumption of the old depends on the capital return. In the natural equilibrium, the young cannot take on any return risk because they enter the market after the current return is realized.

In the complete-markets equilibrium, agents can in principle trade the same history-contingent consumption goods as in our main model. However, since it is not feasible to move consumption across periods, in equilibrium no risk is spread over time. Risk sharing occurs only between the old and young within a period. In effect, the only relevant markets are those for consumption in a period conditional on the state in that period, not the whole history. For the case of small shocks, equilibrium consumption of the old and young are given by

$$c_t^y \approx 1 - q + D_t \left[\frac{\sigma(1 - q)}{\mu + 1} \right],$$

$$c_t^o \approx \mu + q + D_t \left[\frac{\sigma(\mu + q)}{\mu + 1} \right].$$

This solution is the same as the natural equilibrium except that consumption equal to $\sigma(1 - q)/(\mu + 1)$ is shifted from the old to the young

when the capital return is high and vice versa when it is low. The ratio of consumption in the two states is the same for the two generations, as required for optimal risk sharing.

While complete-markets risk sharing is more limited here than in our main model, the broad policy implications are similar. One can show that the main ideas of propositions 3–5 still hold: the natural equilibrium is suboptimal; the complete-markets equilibrium can be implemented by a fully funded social security system that holds capital (in this case the tree, which is expropriated from the first generation); and the social security system can also hold safe debt, allow private agents to hold the tree, and adjust benefits to offset shocks to the tree's output.

Some details of optimal social security are different from before. For example, in our main model, the level of benefits is approximately a random walk when the trust fund holds equity. Taxes also follow a random walk if agents consume in both periods of life. Here, benefits and taxes are serially uncorrelated. A high capital return implies higher benefits and lower taxes in the current period but does not affect the future. This result reflects the impossibility of smoothing consumption over time when there is no investment.

VIII. Conclusion

This paper has explored an approach to analyzing intergenerational risk sharing. According to this approach, policy makers 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 retirement 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 index-

ation 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 4. 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, e.g., Feldstein and Samwick 1999; 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 5. Either approach could implement the complete-markets 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. The model in this paper makes many strong assumptions: individuals within a generation are homogeneous, 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 complete-markets equilibrium—the equilibrium in an overlapping-generations model with complete Arrow-Debreu contingent-claims markets in a Rawlsian original position—is quite general. Future work could investigate the nature of the complete-markets 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 complete-markets equilibrium when there is a single capital-return shock in period j . We do this by deriving two necessary conditions for the equilibrium, equations (4) and (5).

To derive equation (4), we start with the no-arbitrage condition for firms, equation (3). This condition says that the price of a good at t equals the revenue from using the good to produce goods at $t + 1$. In the one-shock case, the price of a good at $t \geq j$ is P_{ti} for $i = H, L$. The revenue from the good is $\mu P_{(t+1)i}$, since each unit of good ti produces μ units of good $(t + 1)i$. (In contrast to the many-shock case, only a single good at $t + 1$ is produced, because the two states at $t + 1$ collapse to one.) Setting $P_{ti} = \mu P_{(t+1)i}$ yields equation (4).

Now consider equation (5), which gives the relative price of consumption in

histories H and L . 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 two histories. With utility function (1), the marginal utility is $1/c$. Thus $Q = c_{tH}/c_{tL}$ for $t \geq h$.

The other condition underlying (5) 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 μ that holds from $j+1$ on. The present value of resources is the gross return on capital at j plus the present value of endowments 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 reverse were true, the allocation would be inefficient and hence could not be a Walrasian equilibrium.

In the high-return history, the gross capital return in period j is $\mu + \sigma_j$ and the endowments at j and subsequent periods are all one. The present value of these resources is $\mu + \sigma_j + [\mu/(\mu - 1)]$. In the low-return history, the endowments are the same but the capital return at j is $\mu - \sigma_j$; the present value of resources is $\mu - \sigma_j + [\mu/(\mu - 1)]$.

Since the present value of resources must equal the present value of consumption in each history, the ratio of the present values of resources in the two histories must equal the ratio of the present values of consumption. We know that the ratio of consumption in the two histories, c_{tH}/c_{tL} , is Q in every period $t \geq j$. Therefore, the ratio of present values of consumption is Q . Setting Q equal to the ratio of present values of resources yields

$$Q = \frac{\mu + \sigma_j + [\mu/(\mu - 1)]}{\mu - \sigma_j + [\mu/(\mu - 1)]},$$

which implies equation (5).

Equilibrium Consumption Levels with a Single Capital-Return Shock

Given the relative prices in equations (4) and (5), one can derive equilibrium consumption levels from agents' utility maximization problems. An agent born in period $t \geq j$ has an endowment of one unit of good tH and one unit of good tL . If we treat good tH as the numeraire, the value of this endowment is $1 + Q$. The agent wants to consume goods $(t+1)H$ and $(t+1)L$. Given (4) and (5), his budget constraint is

$$c_{(t+1)H} + Qc_{(t+1)L} = \mu(1 + Q). \quad (A1)$$

Maximizing expected utility, equation (1), over the two histories subject to (A1) yields the solutions for $c_{(t+1)H}$ and $c_{(t+1)L}$ in equation (7).

The generation born at $j-1$ receives a unit of good $j-1$: this is not indexed by H or L because the shock has not yet occurred. From condition (3), the value of this endowment is the same as the value of $\mu + \sigma_j$ units of good jH and $\mu - \sigma_j$ units of good jL . An agent born at $j-1$ purchases goods jH and jL subject to the budget constraint

$$c_{jH} + Qc_{jL} = \mu + \sigma_j + Q(\mu - \sigma_j). \quad (A2)$$

Maximizing expected utility subject to (A2) leads to the consumption levels in (6).

A Planner's Problem

As discussed in the text, the complete-markets equilibrium is efficient. Therefore, this allocation is the solution to a social planning problem for some set of weights on the utility of different generations. This fact will help us prove propositions 1 and 2.

Formally, the planning problem is

$$\max_{\{c_t(\theta^t), i_t(\theta^t)\}} \sum_{t \geq 1} \beta(t) \left[\sum_{\theta^t \in \Theta \times \dots \times \Theta} u(c_t(\theta^t)) \left(\frac{1}{2}\right)^t \right] \quad (\text{A3})$$

subject to the resource constraint

$$i_t(\theta^t) + c_t(\theta^t) = 1 + i_{t-1}(\theta^t)R_t,$$

given $i_0 = 1$. The constant $\beta(t)$ is the weight the planner puts on the generation that consumes at t . Note that $(\frac{1}{2})^t$ is the probability of each of the 2^t possible histories through t .

The planner's problem produces a familiar first-order condition:

$$u'(c_t(\theta^t)) = \left[\frac{\beta(t+1)}{\beta(t)} \right] \frac{1}{2} [u'(c_{t+1}(\theta^t, H))(\mu + \sigma_{t+1}) + u'(c_{t+1}(\theta^t, L))(\mu - \sigma_{t+1})]. \quad (\text{A4})$$

For any history θ^t , the left side of this condition is the marginal utility of consumption at t . The right side is a ratio of planner weights times the expected value of $u'(c_{t+1})R_{t+1}$, given the two possible continuations of history θ^t .

Proof of Proposition 1

To prove proposition 1, we use the fact that the complete-markets equilibrium solves the planner's problem in (A3). We show that the planner chooses higher consumption at t if there is a high rather than low capital return at $\tau \leq t$.

For this argument, we simplify the planner's problem by noting that there is only one state variable. This variable is the gross return on the previous period's investment, $i_{t-1}R_t$, which we will denote x_t . The planner begins each period with resources $x_t + 1$, so the history θ^t affects his opportunities only through x_t . Therefore, we can write the planner's optimal behavior as a function of x_t rather than the whole history: $c_t(x_t)$ and $i_t(x_t)$.

Given this setup, there are two steps in proving proposition 1. First, we show that $0 < c'_t(x_t) < 1$ for all t . Second, we show that this fact is sufficient to prove the proposition.

We will only sketch the first step since the formal details are tedious. The condition that $0 < c'_t(x_t) < 1$ is equivalent to $0 < i'_t(x_t) < 1$, since $i_t = x_t + 1 - c_t$. The condition says that an extra unit of resources brought into a period goes partly to consumption and partly to investment, which raises future consumption. This is standard consumption smoothing. A windfall must be split between the present and future to maintain the optimal ratios of expected marginal utilities (eq. [A4]).

Proposition 1 states that c_t is higher in history θ^t than in θ'' , which differs only by having a lower capital return at $\tau \leq t$. Since $c'_t(x_t) > 0$, it suffices to prove that x_t is higher in history θ^t : $x_t(\theta^t) > x_t(\theta'')$. This follows by induction if (i) $x_\tau(\theta^t) > x_\tau(\theta'')$ and (ii) $x_s(\theta^t) > x_s(\theta'')$ implies $x_{s+1}(\theta^t) > x_{s+1}(\theta'')$ for $\tau \leq s < t$. Fact i is true because $R_\tau(\theta^t) > R_\tau(\theta'')$ and $i_{\tau-1}(\theta^t) = i_{\tau-1}(\theta'')$, since the two histories are identical

through $\tau - 1$. To establish fact ii, note that $x_s(\theta') > x_s(\theta'')$ implies $i_s(\theta') > i_s(\theta'')$ because $i'_s(x_s) > 0$. By assumption, $R_{s+1}(\theta') = R_{s+1}(\theta'')$, so $x_{s+1}(\theta') > x_{s+1}(\theta'')$.

The Case of Many Small Shocks

Equation (8) gives a first-order approximation in σ_j of equations (6) and (7), the equilibrium consumption levels in the example of a single capital-return shock. To see how (8) is derived, consider c_{tH} for $t \geq j + 1$. Evaluating the expression in (7) at $\sigma_j = 0$ yields $c_{tH} = \mu$ since $Q = 1$ when $\sigma_j = 0$. Differentiating c_{tH} with respect to σ_j yields

$$\begin{aligned} \left(\frac{dc_{tH}}{d\sigma_j} \right) \Big|_{\sigma_j=0} &= \left[\left(\frac{dc_{tH}}{dQ} \right) \Big|_{Q=1} \right] \left[\left(\frac{dQ}{d\sigma_j} \right) \Big|_{\sigma_j=0} \right] \\ &= \left(\frac{\mu}{2} \right) \left[\frac{2(\mu - 1)}{\mu^2} \right] \\ &= \frac{\mu - 1}{\mu}, \end{aligned} \tag{A5}$$

where the second line uses (7) and (5). This result leads to the approximate solution for c_{tH} in (8). The results for c_{tL} 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 case of $\sigma_t > 0$ for all $t \leq T$. We will focus on deriving consumption in periods $t \leq T$ (the analysis for $t > T$ is similar). We are interested in finding $c_t(\theta'; \sigma_1, \dots, \sigma_T)$, that is, consumption in period t and history θ' as a function of the sizes of shocks. A first-order approximation in $\sigma_1, \dots, \sigma_T$ yields

$$c_t \approx c_t(0, \dots, 0) + \sum_{s=1}^T \frac{\partial c_t}{\partial \sigma_s}(0, \dots, 0) \sigma_s, \tag{A6}$$

where we suppress the θ' .

In this expression, the first term on the right is μ , the equilibrium consumption level when there are no shocks. Within the sum, a term $\partial c_t / \partial \sigma_s(0, \dots, 0)$ can be determined as follows. Consider $c_t(0, \dots, 0, \sigma_s, 0, \dots, 0)$, that is, c_t as a function of σ_s when all other σ 's are zero. This function is the solution for conditional consumption in period t when there is a single capital-return shock at s . Differentiating this solution with respect to σ_s and evaluating it at $\sigma_s = 0$ yields $\partial c_t / \partial \sigma_s(0, \dots, 0)$.

For $s \leq t$, $c_t(0, \dots, 0, \sigma_s, 0, \dots, 0)$ is given by equations (6) and (7). We have seen that the derivatives of (6) and (7) evaluated at $\sigma_s = 0$ are $(\mu - 1)/\mu$ when the capital return at s is high and $-(\mu - 1)/\mu$ when the return is low. These results imply $\partial c_t / \partial \sigma_s(0, \dots, 0) = D_s(\mu - 1)/\mu$ for $s \leq t$. For $s > t$, $c_t(0, \dots, 0, \sigma_s, 0, \dots, 0) = \mu$ and $\partial c_t / \partial \sigma_s(0, \dots, 0) = 0$. Substituting these results into (A6) yields equation (9), the approximate solution for consumption in the case of many shocks.

References

- Abel, Andrew B., N. Gregory Mankiw, Lawrence H. Summers, and Richard J. Zeckhauser. 1989. "Assessing Dynamic Efficiency: Theory and Evidence." *Rev. Econ. Studies* 56 (January): 1–19.
- Ball, Laurence, and N. Gregory Mankiw. 2001. "Intergenerational Risk Sharing in the Spirit of Arrow, Debreu, and Rawls, with Applications to Social Security Design." Working Paper no. 8270 (May), NBER, Cambridge, MA.
- Bohn, Henning. 1998. "Risk Sharing in a Stochastic Overlapping Generations Economy." Manuscript (December), Dept. Econ., Univ. California, Santa Barbara.
- Diamond, Peter A. 1965. "National Debt in a Neoclassical Growth Model." *A.E.R.* 55 (December): 1126–50.
- Feldstein, Martin, and Elena Rangelova. 2001. "Individual Risk in an Investment-Based Social Security System." *A.E.R.* 91 (September): 1116–25.
- Feldstein, Martin, and Andrew Samwick. 1999. "Maintaining Social Security Benefits and Tax Rates through Personal Retirement Accounts: An Update Based on the 1998 Social Security Trustees Report." Working Paper no. 6540 (March), NBER, Cambridge, MA.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence." *J.P.E.* 86 (December): 971–87.
- Lucas, Robert E., Jr. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (November): 1429–45.
- Rangel, Antonio, and Richard Zeckhauser. 2001. "Can Market and Voting Institutions Generate Optimal Intergenerational Risk Sharing?" In *Risk Aspects of Investment-Based Social Security Reform*, edited by John Y. Campbell and Martin Feldstein. Chicago: Univ. Chicago Press (for NBER).
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard Univ. Press.
- Shiller, Robert J. 1999. "Social Security and Institutions for Intergenerational, Intragenerational, and International Risk-Sharing." *Carnegie-Rochester Conf. Ser. Public Policy* 50 (June): 165–204.
- Smetters, Kent A. 2002. "Trading with the Unborn: A New Perspective on Capital Income Taxation." Manuscript (November), Wharton School, Univ. Pennsylvania.