University of Pennsylvania
ScholarlyCommons

Wharton Pension Research Council Working
Papers

Wharton Pension Research Council

1-1-2005

Who Bears What Risk? An Intergenerational Perspective

Henning Bohn

University of California Santa Barbara, bohn@econ.ucsb.edu

Follow this and additional works at: https://repository.upenn.edu/prc_papers

Part of the [Economics Commons](#)

Bohn, Henning, "Who Bears What Risk? An Intergenerational Perspective" (2005). *Wharton Pension Research Council Working Papers*. 380.

https://repository.upenn.edu/prc_papers/380

The published version of this Working Paper may be found in the 2006 publication: *Restructuring Retirement Risks*.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/prc_papers/380

For more information, please contact repository@pobox.upenn.edu.

Who Bears What Risk? An Intergenerational Perspective

Abstract

Many governments promise pension and medical benefits to their elderly citizens. As the world is aging, the burden of retiree benefits is becoming painfully obvious. Uncertainty about the future makes planning for retiree benefits even more difficult. Who will suffer or gain financially if the future differs from what we expect? For example, we face tremendous uncertainty about the speed of technical progress, about medical costs, and about trends in fertility and longevity. Government policy determines not only the level of taxes and benefits, but also who bears the risk of unexpected changes. Traditionally, retirement programs largely exempted retirees from sharing risk: by making fixed, unconditional, promises, they necessarily imposed a more-than-proportional risk on younger cohorts and on future generations. We examine the impact of alternative tax, pension, and health care policies on different cohorts, to evaluate how existing policies shift risk across cohorts. We also assess whether there may be conditions under which such policies might be appropriate in the interest of general welfare, and where there may be scope for better policies. The analysis covers the fundamental sources of risk: productivity, fertility, longevity, health, and assets

Disciplines

Economics

Comments

The published version of this Working Paper may be found in the 2006 publication: [Restructuring Retirement Risks](#).

Chapter 2

Who Bears What Risk? An Intergenerational Perspective

Henning Bohn

Governments in most developed countries promise pension and medical benefits to their elderly citizens. As the number of retirees is growing rapidly, the burden of retiree benefits has become painfully obvious. Uncertainty about the future complicates the planning for retiree benefits. If the future is brighter than expected, who will reap the gains? If it is worse, who will cover the added cost?

This chapter examines the impact of uncertainty on different cohorts at an aggregate or macroeconomic level. Macroeconomic analysis helps evaluate different risks and the policies affecting them in context, and it also enforces a consistent recognition that society as a whole faces an uncertain future and must bear the resulting financial risks. These risks include a tremendous uncertainty about the future path of technical progress, medical innovations, and trends in fertility and longevity—risks so huge that common stock market risks are small in comparison. Tax, pension, and health care policies have a major impact on who bears these risks.

Risk-sharing is instructive as a general perspective on public policy because it avoids divisive battles about redistribution. Everyone is better off if risks are shared. Risk should be seen as a symmetric chance of outcomes better or worse than expected. Furthermore, economic risks are often compensated by gains in expectation, creating interesting risk-return trade-offs. The key challenge for economic policy is therefore not to minimize risks, but to allocate risks to those best able to bear them. If risk-taking is rewarded in the market, a related challenge is to maintain incentives for risk-taking and to focus policy interventions into areas where markets fail. Once actual outcomes are observed, however, public policy is inevitably a battle between known winners and known losers, leaving little scope for disinterested economic analysis. Risk-sharing is not only a natural perspective looking forward—a search for mutually beneficial insurance arrangements—but equally instructive looking back: to what extent can existing social institutions be explained as solving risk-sharing problems?

Intergenerational risk-sharing is fertile ground for finding market failures because future generations are naturally excluded from insurance

markets. Welfare improvements are possible because a government's power of taxation gives it a unique ability to make commitments on behalf of future generations. Fiscal institutions such as social security and Medicare formalize such commitments. The government's power to oblige future generations also creates potential for abuse. Risks might be shifted haphazardly onto future generations by governments catering to current voters. The merits of government intervention are therefore an open question—a question inviting economic analysis.

In this chapter, we first identify the key issues and mechanisms of risk-sharing and illustrate them with policy examples. Next we examine three major risk factors: (a) macroeconomic risks, particularly uncertain productivity growth and uncertain asset values; (b) demographic risks, due to uncertain fertility and longevity; and (c) medical expense risks, due to uncertain health care needs and cost. For each risk factor, allocations of risk under current and proposed policies are compared to efficient risk-sharing and to a *laissez-faire* allocation.

Aggregate and Generational Risks

The future rarely unfolds as expected. Over the typical life cycle, individuals face uncertainty about earnings and job prospects, their health and family status, the return on their savings, and ultimately about the time and manner of their deaths. Some of these risks imply financial burdens that add up across cohorts and over time—driven by macroeconomic disturbances—while other risks wash out.

On a macroeconomic level, uncertainty about earnings and the returns on saving have common roots, namely uncertainty about technological progress that determines factor productivity and asset values, as well as uncertainty about demographic developments that determine the supply of labor relative to capital. In turn, demographic uncertainty can be attributed to more fundamental shifts in fertility, longevity, and health (e.g. disability). Health and mortality can perhaps be traced to even more fundamental factors such as innovations in medical technology.

Tracing risks to their fundamental sources is most easily done at an aggregate level to avoid the confounding effects of idiosyncratic noise. Tracing risks is important for economic and policy analysis because common sources of risk create positive correlations that make risk-sharing difficult. Because individuals are exposed to different sources of risk as they age, it is instructive to aggregate risks by cohort (birth year) or by generation (a collection of cohorts). Risks that remain significant for an entire generation are essentially macroeconomic. Such risks are more difficult to manage than idiosyncratic risks because they may not cancel out. Risk-sharing is nonetheless promising because different generations are often exposed unequally to the various sources of aggregate risk.

Managing Aggregate Risks

Three main mechanisms exist for managing aggregate risks: markets, families, and governmental fiscal policy. Financial markets and insurance markets have serious limitations with regard to aggregate risks. In practice, markets work well for sharing many short-term risks between cohorts with largely overlapping lifetimes. However, they cannot provide insurance when there is a very large age difference between cohorts, one large enough so that the older cohort's life risks are largely known, when the younger cohort arrives. In aggregate, the insurance industry is 'owned' by the same generation of savers for which the industry provides insurance. Similarly, corporate pension promises are made by firms that are collectively owned by cohorts as they approach retirement. Consequently, neither insurance policies nor private pensions can provide significant protection against aggregate risks.

The most promising venue for private risk-sharing is probably intergenerational risk-sharing, at least, if risks are defined at the national level. Intergenerational risk-sharing has been puzzlingly ineffective in practice, however. Saving has historically flowed mostly into domestic investment. Investment portfolios are strongly biased toward domestic securities. The analysis later will therefore take mostly a closed-economy perspective, as appropriate in a world with strong home bias.

A second mechanism is intrafamily exchange and altruism, as expressed through bequests and *inter vivo* gifts. In theory, altruism can fully solve all risk-sharing problems, especially if combined with an intracohort sharing of idiosyncratic risk. As shown by Barro (1974), the Ricardian neutrality proposition has dynastic families behaving like a single infinitely lived economic agent, making generational issues moot.¹ Empirically, however, risk-sharing within families is highly imperfect, as documented by Altonji et al. (1996). Parental altruism toward children is undoubtedly an important explanation for education expenditures, but overall the life-cycle model is a good first approximation for individual behavior. Risk-shifting through bequests may nonetheless occur 'accidentally', if mortality is uncertain and assets are not annuitized.

The third risk-sharing mechanism is fiscal policy, via social insurance programs, general taxes and transfers, and through public debt. The most important international risk-sharing programs in the USA are social security old-age benefits and retiree medical insurance (known as Medicare and Medicaid). Public debt also plays an equivalent role for intergenerational redistribution, as it tends to be refinanced repeatedly and passed on to future generation just like pay-as-you-go (PAYGO) pensions. In this setting, taxes are a general purpose risk-sharing device: they socialize the tax share of whatever tax base they are imposed on. Particularly important in the generational context are consumption taxes and capital income taxes; the

latter broadly construed to include taxes on saving, including individual interest, dividend, and capital gains income, corporate income, and on real property.

Economists have shown that fiscal risk-sharing is sometimes constrained by tax distortions, as the excess tax burden grows more than proportionally as labor income tax rates rise. Because risk-sharing leaves taxes rates unchanged on average, distortionary taxes per se are not an issue. Also not at issue are capital income taxes: if stochastic capital income is to be taxed for risk-sharing purposes, saving distortions can be avoided by compensating up-front incentives. A convex excess burden does imply a welfare loss, if risk-sharing calls for variations in labor income tax rates. This is conceptually an enforcement problem inherent in all insurance. After an insured event occurs, one or the other party must be forced to pay up. For intergenerational risk-sharing, convexity means that payments are more costly to collect from future generations than they are to disburse. The welfare impact is roughly proportional to the labor supply elasticity and could be minimized by collecting risk-sharing related taxes at times in the life cycles where labor supply elasticities are small. The distortion issue is noted later where relevant.

Three Policy Examples

Examples are instructive to illustrate the role of policy and to convey two key points. First, none of the risk-sharing mechanisms can eliminate the underlying risks. Therefore, one should be skeptical of policies offering safety to some group, without disclosing who is supposed to provide the guarantees. Second, details matter. Risk-bearing is often determined by subtle features of economic institutions that are often ill-defined or poorly understood.

Who Ensures That Social Security Is Safe? Social security—by which we mean old-age benefit programs—are known to experience recurrent financial problems, which are mainly attributable to the contradictory way in which such programs account for risk. Benefits tend to be set—in the defined benefit (DB) world at least—according to a fixed formula. In the USA, payroll tax rates are also supposed to be constant, but even if budgets initially balance, the risk profiles of benefits and taxes are inconsistent. Consequently, benefits and taxes are bound to drift out of balance if the economy does not grow exactly as expected, or if demographic trends fail to exactly match expectations. Honest planning should acknowledge that either tax rates or benefits must vary in response to economic and demographic disturbances.

In the reform milieu, two paradigmatic structures can be envisioned: a DB model, which will inevitably entail variable taxes to pay the set benefits

14 Henning Bohn

or a Defined Contribution (DC) model, which entails variable benefits. DB and DC are instructive points of reference (used later) because more complicated transfer systems are conveniently characterized by how they compare to DB and DC.

For the US social security system, which is a DB model, having a Trust Fund has complicated an assessment of the system's risk-sharing structure. Having the Trust Fund does help smooth temporary fluctuations in revenues and benefits, but it cannot solve the deeper problem of uncertainty. In 1983, Alan Greenspan's Commission to reform social security raised taxes and reduced benefit growth rates, apparently believing that building up a surplus in the Trust Fund would fix the system's uncertainties. Nevertheless, even if true in expectation, the fix was bound to fail in practice. Since 1983, the Trust Fund and the resulting chaos in federal accounting (Trust Fund accounting conflicts with national Unified Budgeting) have obscured the fact that policymakers have not addressed the key questions of whether social security should be seen as a DB or a DC system, or something in between.

If social security is viewed as a DB model, then it is reasonable to think of it as creating 'safe' claims; in other words, it entails an unconditional obligation on future generations to finance promised benefits. To honor this obligation, tax rates will have to rise whenever payrolls grow less than expected, and vice versa. By contrast, if it were seen as a DC model, social security benefits would be contingent claims worth as much (or as little) as the payroll tax revenue dedicated to them. Either way, future benefits payable and taxes collected are influenced by a multitude of economic and demographic disturbances.

The DB versus DC question is crucial for those who expect to live off social security in old age, but there is little agreement over which is the right interpretation. US political rhetoric for many years supported the DB view, but more recently the government has adopted more of a DC perspective, as seen in social security statements warning of benefit cuts when the Trust Fund is exhausted. If widely accepted, this new view would have a major impact on risk-sharing in the USA, as it would expose retirees to additional risk of future economic and demographic changes. If increases in payroll taxes are deemed to be politically infeasible, this supports the DC interpretation. We examine both interpretations in later subsections.

What Are the Risk-Sharing Implications of Social Security Reform? The policy debate over whether to hold equity investments in social security illustrates some of the risk-sharing implications of seemingly minor differences in program design (Bohn 1999). For instance, consider two alternative proposals: permitting workers to convert part of their social security taxes to a DC-type individual account, as suggested by the 2001 Presidential Commission, versus the Clinton administration's plan, which

proposed to invest part of the Trust Fund in the stock market while maintaining a DB promise. The individual accounts approach reduces risk-sharing, scaling down the traditional system proportionally. Under the Clinton plan, by contrast, equities in the Trust Fund are effectively owned by future tax payers, so equity risk is shifted across generations.

Now consider two modifications that might be thought of as relatively modest. First, suppose the individual accounts were combined with a minimum-return guarantee. Financial economists know that a return guarantee is economically equivalent to a put option on the stock market, which in turn is equivalent to a transfer of fractional ownership to future generations which back the guarantee (Smetters 2001). In this sense, an individual account plan with a guarantee has remarkably similar risk-sharing implications to the Clinton plan. Furthermore, the put option argument applies even without explicit guarantees, if those who earn low returns in their individual accounts would be eligible for welfare benefits.

As a second example, consider the Clinton approach to equity investments in the Trust Fund, but now with a DC view of social security where gains and losses in the Trust Fund would accrue to retirees. Clearly, the risk-sharing implications are equivalent to nonguaranteed individual accounts: no intergenerational risk-shifting. Trust Fund investments rule out individual debt-equity choices, of course. Even this microeconomic difference vanishes if individual accounts are restricted to a single index fund.

In summary, the risk-sharing implications of equity investment proposals depend on the specifics. Individual account plans imply very little effective ownership if accompanied by guarantees or by means-tested supplemental supports. Conversely, having the Trust Fund diversify its investments has very different risk consequences, depending on whether the system is seen as a DB versus a DC plan.

Who Bears the Financial Risks of Medicare? In the USA, retiree medical benefit promises have traditionally been presented as a categorical commitment, just like social security. Retiree Hospital Insurance (HI, known as Medicare Part A) is financed by the same payroll-tax mechanism as are old-age benefits. Medicare is of interest because of two important complications: substantial uncertainty about medical expenses and the cost-sharing structure embedded in the Supplemental Medical Insurance plan (SMI, known as Medicare Part B and Part D). If Medicare benefits are seen to be a promise to ‘cover all medical needs’, then it follows that the system can be interpreted as a DB plan with less predictable costs than the old-age social security system. In this view, benefits are contingent on the uncertain cost of existing treatments and contingent on medical discoveries that may vastly increase beneficiaries’ perceived needs. Both create what I will call *medical-expense risk*. Cost projections under this DB interpretation are enormous and

16 Henning Bohn

highly sensitive to alternative assumptions. Clearly, this interpretation exposes future generations to substantial risks, to be examined in more detail later.

Cost-sharing in SMI complicates the assessment, as SMI has always required retiree contributions, and there is a history of cost-shifting between HI and SMI. To the extent that rising health costs could be shifted to SMI, retirees will retain some of the aggregate medical-expense risk. On the other hand, most SMI coverage is financed from general tax revenues, and many retirees have limited income. As long as medical care in retirement is viewed as entitlement, medical-expense risk is thus largely carried by future generations.

This allocation of risk would be reversed completely if aggregate expenditures were capped. For example, Kotlikoff and Burns (2004) propose that Medicare be converted into a voucher program subsidizing individual insurance purchases. They argue, somewhat misleadingly, that vouchers provide full insurance and are sensitive to preexisting conditions. But a key assumption in their approach is that vouchers increase in value no faster than the growth rate of wages; if aggregate medical expenses rose faster, the vouchers would cover a declining share of each retiree's medical expenses. A capped voucher system therefore places aggregate medical-expense risk onto the retiree generation.

These examples illustrate how program rules determine the allocation of risk within and across generations. The lesson is clear—there is no free lunch. With proper accounting, governments cannot make macroeconomic risks vanish; instead, the risk is simply reallocated. The observation leads to the next question: how *should* society allocate aggregate risks?

A Benchmark: Equal Risk-Sharing

A natural benchmark for risk-sharing is the proportionate pooling of risk. Everyone bears the impact of any conceivable good or bad economic outcome in proportion to his normal consumption opportunities. Risk-pooling is the unique efficient allocation if everyone has the same relative risk aversion. If risks were allocated differently, individuals with above-average risk exposure would be willing to pay a higher price for risk-reduction than individuals with below-average exposure. Unequal risk exposures are therefore indicators of economic inefficiency.

Risk-pooling provides a clear guideline for policy and a roadmap for economic analysis. To improve welfare, policy should shift risks from cohorts more exposed to a given source of risk to those initially less exposed. To find policy improvements, one must assess who bears how much of each risk in the market, how the risks would be allocated with efficient sharing, and then compare the market allocation with the efficient allocation.

Two clarifications and two caveats are in order. First, efficient policy regarding risk-bearing and risk-sharing should focus on consumption (or more precisely, the marginal utility thereof) and not income. Because labor supply and savings opportunities differ across cohorts, consumption pooling usually entails unequal exposures to income risk. Second, risks can be shared with future generations by variations in capital accumulation. Most economic disturbances therefore trigger consumption responses in current and future periods. Efficiency requires that retiree consumption and working-age consumption respond equally in every period, usually with declining amplitude over time.

The first caveat is to emphasize that aggregate risks have a market price. In finance terms, aggregate risks are systematic risks. On financial markets, securities with systematic risk trade at a discount to safe assets, or equivalently, they promise an expected return above the safe interest rate. The same principle applies to aggregate risks that are reallocated through fiscal policy and may be nontradable. To make everyone better off, a policy must offer compensation to those required to bear more risk—a risk premium. The acid test for a welfare-improving policy is that those relieved of risks are better off *after* paying the risk premium. Unfortunately, economic theory has trouble explaining empirically observed risk premiums, notably the equity premium. While risk pooling provides a straightforward benchmark for how risk should be allocated, the magnitude of the compensation is sometimes difficult to determine.

The second caveat is that risk exposures must be adjusted if individuals differ in their intrinsic risk aversion. Less risk-averse individuals demand a lower price for bearing risk and should bear more risk in an efficient allocation. The literature on habit formation suggests that older people, the retiree cohort, may be more risk-averse than younger, working-age individuals. Later I will return to this caveat.

In any case, efficient risk-sharing has implications that are robust to differential risk aversion. Most importantly, for economic efficiency, everyone would have to be exposed to aggregate risk in the same direction. If a disturbance hit only working-age cohorts, or only retirees, or both generations in opposite directions, risk-sharing is always inefficient. Moreover, if one cohort is more risk-averse than another, it should be proportionately less exposed to all types of risk, yielding testable restrictions in a world with many sources of risk.

A Tractable Analytic Framework

The risk-pooling principle has most power if applied across multiple risks. Because examining and aggregating the joint effects of all the various risks on different cohorts is not an easy task, it is instructive to take a stylized perspective on work and asset accumulation over the life cycle to see how

the principle can be evaluated. The classic Diamond (1965) two-period overlapping-generations (OG) model provides a convenient framework.

In the Diamond model, all cohorts in the workforce are grouped together into a single working-age generation, and all retirees are treated as a single retiree generation. Children are attached to their parents' households. Changes in longevity can be modeled as changes in the relative length of working-age and retirement. While this description of the life cycle is clearly simplified, it helps highlight key elements of intergenerational risk-sharing without getting lost in the minutiae of how individuals lead their lives. Incomes during the working-age period are mainly wage incomes; wages after taxes are divided between consumption and saving. Saving is invested in financial markets, and it ends up financing either domestic capital accumulation, government bonds, or investments abroad. Income for the retiree generation is from capital and other assets. Retirees use asset income, asset sales, and transfers from the government to finance retirement consumption, medical care, and (perhaps) bequests.²

What can this framework tell us about who bears what risks? The following sections will examine each of the major risks.

Aggregate Risks (I): Macroeconomic Risks

Uncertain Productivity Growth. Productivity growth is quantitatively the most important source of long-run economic uncertainty. With 1 percent annual productivity growth for a generation, for example, our children will earn 35 percent higher incomes (compounding 1 percent for 30 years), and our grandchildren 80 percent more. With 3 percent annual productivity growth, our children would instead earn 140 percent more and our grandchildren almost 500 percent more. Without productivity growth, per capita incomes would stagnate. The 0–3 percent range falls well within the range growth rates experienced around the world.

Who bears the risk of such growth uncertainty? Because labor and capital shares in national income are essentially constant, growth uncertainty has an equal impact on wages earned by working-age households and on capital incomes earned by retirees. But equal income effects do not imply equal consumption effects; there are three main issues to consider. First, the retiree generation owns not only the earnings but also the principal value of accumulated real and financial assets. Because asset values are less sensitive to growth than earnings (see later), existing assets reduce retiree generation's exposure to productivity risk, leaving workers relatively more exposed. Second, government transfers and taxes augment retiree incomes at the expense of working cohorts' incomes. Because net transfers tend to be less sensitive to growth than wages and capital incomes, fiscal policy further reduces retiree exposure to productivity risk while increasing working cohorts' exposure. Third, the working-age generation may respond to

unexpectedly high or low productivity by altering its saving rate and work hours. The study of labor supply and savings responses is unfortunately complicated because several income and substitution effects interact. In Bohn (2004), I show that savings and labor supply responses tend to further magnify the impact of productivity shock on working-age consumption.

Key distinctions must be made between income and productivity shocks and between temporary and permanent shocks. Negative *income* shocks can be absorbed by working more; however, negative *productivity* shocks reduce not only income but also the hourly wage, which makes a productivity slump an inefficient time to work more. Having variable labor supply therefore does not enable working cohorts to bear more productivity risk. The permanence of shocks matters because the young can absorb temporary shocks by consumption-smoothing over the life cycle. A change in productivity *growth* is, however, a permanent disturbance to the economy's productivity level, which rules out consumption smoothing.³ The permanence of shocks also matters because high future productivity raises interest rates, which has income and substitution effects. Assuming a low elasticity of substitution (consistent with empirical estimates), high-productivity growth reduces the savings rate and therefore magnifies the impact of productivity shocks on working-age consumption.

Overall, the OG approach yields an unambiguous conclusion: *working-age individuals are more exposed to productivity risk than retirees*. This result may seem to conflict with the notion that capital incomes are more risky than wage incomes. There is no conflict, however, if one properly distinguishes productivity risk from asset valuation risk—to be examined next.

Uncertain Asset Values. Returns to capital are subject to a second source of risk, namely the uncertainty about the price at which seasoned capital can be sold to the next generation. This topic includes the much-discussed stock market and housing market risks. Because asset valuation risk falls primarily on the retiree generation, retiree incomes can be more volatile than working-age incomes.

The main social mechanisms for sharing such valuation risk include capital income taxation and bequests. Capital income taxes yield return-contingent government receipts that reduce the need for taxes on future generations. This gives future generations an exposure to the current return on capital. Income taxes unfortunately reduce saving incentives. One policy response in the USA has been to provide savings incentives up-front and fully tax the returns, as exemplified in 401(k) and 403(b) plans and traditional IRAs. Another policy response has been to tax capital gains and dividends at reduced rate, or to exempt capital income entirely (e.g. Roth IRAs). The distinction is important because risks are shared in proportion to the marginal tax rate. A tax system with taxed returns and up-front incentives provides much more risk-sharing than tax-exempt savings.⁴ The other important mechanism is

20 Henning Bohn

bequests. For risk-sharing purposes, bequests are equivalent to a self-imposed capital income tax that benefits future generations.

Evidence on Productivity and Valuation Risks. Empirical evidence is consistent with wage income being relatively more exposed to productivity risk than capital income. I have computed 30-year-ahead ‘generational’ covariance matrices from annual 1875–2002 US gross domestic product (GDP) data and S&P500 prices and dividends (Bohn 2004). Here I find a generational standard deviation for GDP of $\sigma_Y = 35$ percent, a 64 percent standard deviation for equity returns, and a 41 percent correlation between the 30-year-ahead forecast errors.⁵ If one interprets the data in terms of a Cobb–Douglas aggregate production function, the standard deviations of wages and productivity equal the standard deviation of GDP. If capital is financed with equity and debt (about 26 percent debt) and debt is essentially safe, the standard deviation of capital (R^k) can be determined to be $\sigma_{R^k} = (1 - 0.26) \times 0.64 = 47$ percent. One may also decompose the return on capital into a GDP factor, plus an orthogonal asset valuation factor (V), $R^k = \pi \times Y + V$. The above estimates imply a factor loading of $\hat{\pi} = \text{cov}(Y, R^k) / \text{var}(Y) = 0.58$ and $\sigma_V = \sqrt{\sigma_{R^k}^2 - \hat{\pi}^2 \sigma_Y^2} = 43$ percent.

Even my long data-set covers only a small number of nonoverlapping generational periods. The generational variances and covariances are therefore best viewed as point estimates subject to substantial specification uncertainty.⁶ It is nonetheless reassuring that the estimates match economic theory quite well. Notably:

- (1) The factor loading below one indicates that capital returns are less exposed to productivity risk than output and wages—about 58 percent as much.
- (2) The inequality $\sigma_{R^k} > \sigma_Y$ confirms the conventional wisdom that returns on capital investment are more volatile than wages.

Note that these estimates are also consistent with productivity uncertainty as the dominant source of consumption risk at long horizons, especially looking several generations ahead. This is because future generations bear valuation risk for a limited number of years, between work and retirement, whereas their exposure to productivity risk grows with the forecast horizon.

The differences in riskiness between retiree and working-age incomes are promising for intergenerational risk-sharing. They suggest that working-age cohorts benefit from shifting productivity risk to retirees, while retirees benefit from shifting asset pricing risks to subsequent generations.

Linking Income to Consumption: A Quantitative Analysis. Quantitative insights about the efficiency of risk-sharing can be obtained by combining the data with a calibrated version of the OG model. Table 2-1, which builds

2 / Who Bears What Risk? Intergenerational Perspective 21

TABLE 2-1 Productivity and Asset Pricing Risks

	<i>Scenarios (Responses to unexpected changes in percent)</i>					
	<i>Laissez-faire</i>	<i>Laissez-faire with bequests</i>	<i>Efficient solution</i>	<i>Calibrated policy</i>	<i>More taxes and wage-indexing</i>	<i>All safe transfers</i>
	(1)	(2)	(3)	(4)	(5)	(6)
Panel 1A: Higher productivity (35 %)						
Consumption in retirement	26	26	34	25	28	17
Consumption in working age	36	35	34	37	35	41
Disp. Income in working age	35	34	30	35	33	39
Capital investment	32	31	23	31	29	37
Panel 1B: Higher asset values (40 %)						
Consumption in retirement	40	40	13	27	21	27
Consumption in working age	0	4	13	7	10	7
Disp. Income in working age	0	4	16	8	11	8
Capital investment	0	5	22	10	14	10

Source: Author's calculations.

Notes: Columns refer to different policy scenarios, as follows.

Col. 1: Basic overlapping generations economy without government and no bequests.

Col. 2: Like Col. 1 with accidental bequests.

Col. 3: Responses with efficient risk-sharing. (Key characteristic: Equal consumption responses.)

Col. 4: Economy calibrated to US data and fiscal institutions: partially wage-indexed social security, defined-benefits public health, safe public debt, capital income taxes with 15 percent marginal rate.

Col. 5: All transfers wage-indexed, public debt wage- or GDP-indexed, capital income tax with 30 percent marginal rate, otherwise like Col. 4.

Col. 6: All transfers inflation-indexed, otherwise like Col. 4.

on Bohn (2001, 2004), shows how differential income risks translate into unequal consumption exposures.⁷ Each column presents the risk-sharing implications of a different policy scenario. Entries in Panel A show percentage responses in both generations' consumption, in working-age income, and in capital investment to an unexpected 35 percent increase in productivity growth. The 35 percent change can be interpreted as generational standard deviation or as 1 percent per year compounded for

22 Henning Bohn

30 years. Panel B shows responses to an unexpected 40 percent increase in asset values, also about a standard deviation.

The responses are symmetric and scalable. Negative disturbances would have the reverse effects; smaller or greater shocks would have proportionally smaller or greater impact. The percentage values can therefore be interpreted as exposures to risk. They can be used contemporaneously to assess the efficiency of risk-sharing between current retirees and current workers, or prospectively to examine risk-sharing between future retirees (current workers) and future workers (current children and unborn).

Column 1 shows the allocation of risk in a laissez-faire economy without bequests. Higher productivity raises working-age consumption, incomes, and capital investment by about as much as productivity (35 percent), but retirement consumption by substantially less (26 percent). Unexpectedly high asset values affect only retirees. Column 2 adds bequests to the laissez-faire economy, specifically that 20 percent of retiree resources are bequeathed to working-age cohorts.⁸ Retirees remain exposed to the same risks as before, but working-age cohorts bear a little valuation risk. The more diversified income slightly reduces working-age exposure to productivity risk. Column 3 shows the allocation one would obtain with perfectly pooled risk. Retiree and working-age consumption would respond equally to both shocks. Because savings respond more to productivity shocks than working-age consumption, disposable income in working age should be less sensitive to productivity shocks than retiree income. The efficient allocation of valuation risk cuts retirees exposure by more than two-thirds, and it shares valuation risk with future generations through substantial variations in investment. Column 4 adds a stylized representation of US fiscal institutions to the market economy with bequests. The retiree generation receives partially wage-indexed, annuitized social security benefits: 5 percent of GDP on average, 50 percent wage-indexed to proxy indexation to age 60. They receive medical benefits that are unresponsive to productivity risk: 3.5 percent of GDP. They pay income and consumption taxes: 3.5 percent of GDP with 15 percent marginal rate; and they hold safe government bonds: 2.5 percent of generational GDP. The relative magnitudes of transfers and taxes match Auerbach et al. (1999) generational accounts for age 65, rounded for simplicity.

If one compares the calibrated policy (Column 4) with laissez-faire (Column 2) and with the efficient allocation (Column 3), one finds that *fiscal policy magnifies the generational gap between working-age and retiree exposure to productivity risk*. This is largely because public debt, medical benefits, and social security (the nonwage-indexed part) provide safe claims to retirees. While safe transfers reduces retiree exposure to risk, they force the government to collect fixed revenues from the next generation's stochastic wage income. This increases the relative risk exposure of working-age cohorts. Policy also shifts valuation risk from retirees to workers, largely

because of the capital income tax, but not as much as risk-pooling would require.

Columns 5 and 6 illustrate the impact of alternative policies. Column 5 assumes fully wage-indexed transfers: a 30 percent (doubled) tax on capital income, holding savings incentives and the generational account constant through compensating transfers. Column 6 assumes that all retiree transfers are fixed in real terms. The allocation with wage-indexing and higher-marginal tax rates is evidently closer to risk pooling than current policy, whereas fixed benefits yield a more uneven allocation of productivity risk.

Overall, Table 2-1 demonstrates how fiscal policy influences risk-sharing and how it could improve the allocation of risk. Note that a better sharing of productivity risk would make labor income taxes less variable and therefore reduce the excess burden of taxation. Hence tax distortions cannot explain the imperfect sharing of productivity risk. They may help explain why valuation risk is not fully shared.

In general, the public policy debate is quite incoherent regarding how much risk retirees should be expected to take. Two popular positions which embody this conflict are those who would (a) encourage equity investments in pensions and (b) provide guaranteed benefit promises. Both positions are inconsistent with risk pooling. To justify (a), one would have to consider retirees to be highly risk-tolerant, but in this case retirees should bear much more productivity risk, contradicting (b). To justify (b), one would have to view retirees as intrinsically more risk-averse than younger individuals. This might justify fiscal institutions that provide safe income to retirees, but it would argue strongly against pushing retirees into equity investments. If retirees are about as risk-averse as younger cohorts, welfare improvements can be found in the direction of more wage- or GDP-indexing, either within social security or via GDP-indexed public debt.⁹ Not all indexing is beneficial, however. For example, if productivity shocks are negatively correlated with inflation, as evidence suggests, inflation-indexed government debt yields less risk-sharing than traditional nominal debt.¹⁰

Aggregate Risks (II): Demographics

Uncertain Fertility. It is clear that fertility changes can have major effects on the economy and on intergenerational transfers: the baby-boom and baby-bust phenomena are prime examples. Unfortunately, the public pension debate has mischaracterized the phenomenon by focusing on the fiscal burden that falls more heavily on smaller cohorts than on larger ones.

Macroeconomic theory suggests that being born into a small cohort is actually good news. When a small cohort enters the labor force, workers are scarce and retirement saving from the preceding larger cohorts provides a high capital-labor ratio. When small cohorts move toward retirement,

24 Henning Bohn

subsequent larger cohorts provide the labor that allows retirement savings to earn a high return. In a world without PAYGO transfers, risk-sharing would therefore call for net transfers from small cohorts to larger ones. In a world with PAYGO transfers, higher-payroll taxes on smaller cohorts can be interpreted as risk-sharing—not as a fiscal problem, but as solution to a demographic risk sharing problem.

Table 2-2 provides a quantitative assessment of who bears the risk of variable cohort sizes. Panel A displays responses to an unexpected 15 percent decline in the labor force (about the decline between the baby boom and baby bust: 79 million births in 1945–64 versus 69 million births 1965–84).¹¹ Columns 1–6 display the allocation of risk for the same scenarios as in Table 2-1, with the clarification that transfers are of the DB variety in Columns 4 and 5 but DC in Column 6. The laissez-faire allocations

TABLE 2-2 Demographic Risks: Fertility and Longevity

<i>Scenarios (Responses to unexpected changes in percent)</i>						
	<i>Laissez-faire with Efficient bequests</i>			<i>More taxes Calibrated and wage-indexing transfers</i>		
	<i>(1)</i>	<i>(2)</i>	<i>(3)</i>	<i>(4)</i>	<i>(5)</i>	<i>(6)</i>
Panel 2A: Baby-bust decline in the workforce (–15 %)						
Consumption in retirement	–7.9	–7.9	0.1	–4.2	–2.1	–7.6
Consumption in working age	4.0	4.4	0.1	2.2	1.2	3.9
Disp. income in working-age	4.5	4.8	0.1	2.5	1.4	4.4
Capital Investment	5.5	5.6	0.2	3.2	1.7	5.6
Panel 2B: Increased longevity (+15 %)						
Consumption in retirement	–15.0	–15.0	–6.7	–11.6	–8.6	–15.0
Consumption in working age	–5.9	–5.1	–6.7	–6.3	–6.5	–6.0
Disp. income in working age	0.0	0.0	–4.6	–1.9	–3.5	0.0
Capital investment	12.4	12.6	0.0	7.7	3.0	13.0

Source: Author's calculations.

Notes: Same scenarios are as in Table 2-1. Panel A considers a 15 decline in cohort size, roughly matching the US baby-boom to baby-bust transition. Panel B assumes a 15 percent increase in the retirement period (about two years). Transfers are annuitized defined benefits in Cols. 4 and 5, and defined contributions in Col. 6.

illustrate the benefits of being in a small cohort (Columns 1 and 2). The consumption of the retired ‘baby boom’ generation declines, whereas the income and consumption of the working-age ‘baby bust’ generation rise. Efficient risk-sharing calls for a slight increase in both generations’ consumption (Column 3), as permitted by a slight increase in per capita income. The policy (Column 4) calibrated to resemble the US yields consumption responses in the same opposing directions as *laissez-faire*. The smaller absolute values indicate better risk-sharing. The policy alternative with more DB transfers (Column 5) provides more risk-pooling than the calibrated policy, whereas the DC system (Column 6) provides virtually no risk-sharing. In summary, a DB transfer help share demographic risk.

Empirical evidence on the effects of demographic risk is unfortunately scarce. Evidence exists on the labor market, where the wage effects of a variable cohort size are well documented (e.g. Welch 1979). Capital market effects are difficult to document, perhaps because the start of retirement savings is more variable cross-sectionally than a cohort’s entry into the labor force. Capital market effects are also obscured by high asset price volatility and perhaps by international capital flows. Population aging is a worldwide phenomenon, however, so the scope for international diversification is limited. Despite the scarcity of evidence, economic theory is probably a better guide to the future than naive trend extrapolations that assume a disconnect between demographic change and factor prices.

Uncertain Longevity. Around the world, mortality has long been on a declining trend. The main intergenerational impact is on life expectancy in retirement, or longevity, for short. Increased longevity is obviously good news for the retirees. But unless all retirement income is annuitized, increased longevity has a negative impact on living standards. It requires a reduced rate of per-period consumption.

Efficient risk-sharing calls for a sharing of longevity risk with subsequent generations. On a microeconomic level, annuities are the obvious risk-sharing tool, but they are subject to adverse selection and often unavailable. Most assets are held in nonannuitized form, which leads to substantial intergenerational transfers through accidental bequests. On an aggregate level, longevity risk is virtually impossible to insure in the market. Insurance providers are owned collectively by the generation that seeks insurance. Insurance between nearby, but distinct cohorts would have to be signed far in advance, before too much about actual longevity is known. Therefore, the government has a unique role as a provider of longevity insurance backed by future generations, providing a rationale for annuitized public pensions.

The degree of inefficiency without public pensions depends heavily on retiree attitudes toward bequests. To the extent assets are bequeathed intentionally—say, to a close relative or to a favorite charity—there is no inefficiency.

To the extent that bequests occur because annuities are unavailable, public pensions improve retiree welfare. I interpret the popularity of public pensions as indication that most retirees like annuities (i.e. do not have altruistic bequest motives strong enough to make the risk-sharing problem moot).

Not all annuitized public pensions share aggregate longevity risk. The key design issue is if per-period pensions and related medical benefits are fixed regardless of longevity (i.e. of the DB variety), or if payroll taxes and other contributions are held constant (i.e. the DC model). In the former case, increased longevity requires higher taxes. In the latter case, increased aggregate longevity must lead to reduced per-period benefit payments. Table 2-2B provides a quantitative perspective on how fiscal policy allocates longevity risk. It shows responses to a permanent 15 percent increase in longevity during retirement (equivalent to about two years). Under *laissez-faire*, per-period retiree consumption declines one-for-one with longevity (see Columns 1 and 2). Because working-age cohorts also expect a longer retirement, they increase savings and reduce consumption, but only fractionally. Efficient risk-sharing simply calls for sharing current resources (Column 3). Savings remain unchanged because the longevity increase is permanent and leaves no scope for risk-sharing over time. The calibrated allocation (Column 4) shares longevity risk better than *laissez-faire* due to annuitized social security and health benefits, but not completely. An expanded transfer system (Column 5) comes closer to risk-pooling, documenting the risk-sharing role of an annuitized DB model, whereas the DC system (Column 6) is as inefficient as *laissez-faire*.

The retirement age is an open question in this context. One may of course allow early retirement at an actuarially reduced pension. The deeper question is how the normal retirement age should relate to longevity. Medical improvements that extend a generation's ability to work may be viewed as a positive demographic shock. Risk-sharing would then suggest that individuals capable of working longer should indeed work longer and share the gains with other cohorts. The gains may be small, however, because if retirement ages are individually optimal, working longer increases lifetime income by about as much as it reduces utility from leisure, leaving no first-order gains for sharing. A related issue is the baseline for insurance. Only unexpected changes are insurable. If one started with naïve static expectations, any increase in longevity would trigger payments from those living short lives to those who have the financial misfortune of living longer. Insurance against good news—living longer—is counterintuitive from a distributional perspective. If insurance is conditioned on a positive trend path for longevity, in contrast, only deviations from this path are insurable and would trigger increased or reduced intergenerational transfers.

Conditioning on the trend is perhaps the best argument for linking retirement age to longevity. It seems unfair to let generations who die early pay for the predictably higher expenses of longer-living generations.

But to avoid slipping from a DB to a DC world, the normal retirement age needs to be fixed well before a cohort approaches retirement. An exemplary implementation is the increase in retirement age from 65 to 67 in the USA, which was announced decades in advance.

Finally, we note that tax distortions may impose limits on sharing demographic risks. Fertility and longevity risk are shared by imposing higher working-age taxes on unexpectedly small cohorts and on cohorts that follow unusually long-lived ones. It is an open question to what extent tax distortions explain why demographic risk-sharing is incomplete in practice. Distortions do provide another rationale for link normal retirement age to the longevity trend.

Aggregate Risks (III): Health Care

The growth in medical expenses has reached a stage where health care deserves treatment as a macroeconomic risk factor. The social security debate has been remarkably silent about health issues, even though Medicare and Medicaid are growing much more rapidly and—looking forward—will impose the greater fiscal burden. It is tempting to avoid discussing the issue, because health care raises delicate questions about preferences over life, death, and human suffering. I will nonetheless attempt a welfare-theoretic analysis as policymakers simply cannot afford to ignore the issue.

Two very different perspectives are currently influential in the health care debate. One view is that every person is entitled to health and survival, regardless of wealth or income. Under this view, public spending on health care will grow at an exogenous rate, driven by an ethical obligation to pay for all health care that is technically feasible. The economic implications of this are troubling, with fiscal projections totaling \$38 trillion in present value for Medicare (Gokhale and Smetters 2003). Cost uncertainty creates a huge need for insurance. The opposing view is that people buy health care like other commodities; consequently, the demand for health services responds elastically to changing relative prices and to rising income. Growth in medical spending is thus endogenous, driven in part by technological innovations that reduce quality-adjusted prices, in part by rising incomes, and in part by inefficient insurance arrangements that provide false price signals. Insuring expenditures would then be misguided. The insurable events are instead the discoveries of new treatments and the resulting changes in relative prices. The need for insurance depends on health consumers' price elasticity. If innovations reduce the price of curing a health problem, expenditures fall if demand is inelastic; stay constant if demand has unit elasticity, or rise if demand is elastic. In practice, new treatments tend to increase overall medical spending, suggesting that demand is more than unit-elastic. Setting aside the well-known static

inefficiencies of third-party insurance (a topic best left to microeconomics), the key macro implication is that rapid growth in medical spending may well be an efficient response to medical innovations and price-elastic demand.

These contrasting views have some common macroeconomic and intergenerational implications. Medical expenses are bound to increase as share of GDP, either unavoidably (view 1) or because it is optimal (view 2). The generational implications follow from the correlation between age and medical needs: risk-sharing, of whatever type appropriate, will go in the direction of working-age cohorts providing insurance to retirees.

The source of rising medical expenses is important for the distributional baseline. Even if one takes an entitlement view of what society should cover, one must acknowledge that growth in medical expenses is largely driven by new treatments. Intergenerational insurance against health care cost is thus analogous to longevity insurance. Just like longevity insurance, insurance against medical expenses insures against the financial implications of good news—the discovery of new treatments. Just like longevity insurance, such insurance is awkward from a distributional perspective. Why should we be responsible for the medical expenses of future generations who we expect to be more healthy, wealthy, and longer living than we are? The fiscal balance rule of generational accounting appears questionable in this context because it obliges current generations to share the cost of treatments not yet invented (Gokhale and Smetters' \$38 trillion estimate includes some amount of this).

Intergenerational risk-sharing is again about the impact of changes relative to a given baseline, and it is applicable regardless of baseline. For analytic purposes, I divide health care into two conceptually distinct types, roughly corresponding to the two views. Type 1 consists of life-saving or life-lengthening treatments prerequisite for normal life (life-saving care, for short). Type 2 consists of items that make people feel better, here called discretionary care, for short. In the OG model, one may think of life-saving care as affecting the probability of reaching a certain age, but not utility from consumption conditional on survival, whereas Type 2 health care enters into each period's utility function. Risk-sharing with Type 1 medical expenses is centrally about the value of life—an issue where normative answers are outside the scope of economics. If, however, one takes a certain level of life-saving care as given, and if such care enters separably into individual preferences, efficient risk-sharing calls for the unexpected cost to be shared. The usual principles of risk-sharing apply to the generations' nonmedical consumption.

Efficient risk-sharing with Type 2 medical expenses is a straightforward exercise in welfare analysis. Health care prices enter the price index for consumption. Assuming discretionary care has a higher weight in retiree consumption, an unexpected decline in quality-adjusted medical

prices implies a decline in the price index of retiree consumption relative to the price index of working-age consumption. If the elasticity of substitution between health care and other consumption is above one, as presumed in the discretionary view, then efficient risk-sharing calls for retirees to share with working-age cohorts the welfare gains from lower medical prices. Real consumption would increase for all cohorts, but consumption spending would decline for retirees while rising for working-age cohorts.

Table 2-3 provides a quantitative illustration. The setting is the Diamond model described above, now with a distinction between medical and non-medical consumption. Retirees are assumed to have twice as much health care needs as working-age cohorts.¹² Columns 1 and 2 contrast two extreme scenarios, ‘Full Coverage’ of retiree medical expenses (the DB approach) versus ‘Capped Benefits’ (the DC approach). Efficient risk-sharing is shown in Column 3. Consumption spending refers to total consumption, including private- and government-provided health care, with nonmedical consumption as *numeraire*. For reference, public spending in the USA covers about 65 percent of retiree health care. As Medicare drug coverage is rolled out, this percentage is likely to increase.

Panel A explores the implication of a permanent 30 percent increase in medical needs, meaning an increase in expenses that is unavoidable and separable in preferences from other consumption.¹³ The Full Coverage scenario imposes the cost entirely on the next generation of taxpayers, leaving nonmedical retiree consumption unchanged. The Capped Benefits scenario holds overall retiree spending unchanged, forcing them to fund the incremental medical expenses from reduced nonmedical spending. Working-age households also face increased medical expenses, but their response is mainly driven by their expectations about retirement. In the Full Coverage scenario, savings and investment decline as workers pay higher taxes to fund retiree medical expenses without having to worry about increased own expenses in retirement. In the Capped Benefits scenario, working-age savings and investment increase in expectation of higher medical expenses in retirement.

The standard for efficient risk-sharing with separable preferences is an equal response of nonmedical spending. In the setting of Table 2-3A, this implies covering 27 percent of costs (Column 3), forcing retiree and worker nonmedical consumption to decline. Efficiency also implies a zero savings response because expenses increase equally in current and future periods. The 27 percent value is sensitive to the example parameters. The zero investment response is robust—a hallmark of an efficient response to a permanent shock.

Table 2-3B presents a setting where individuals have preferences with constant elasticity of substitution over real medical and nonmedical consumption. Prices are assumed to decline by 30 percent, the elasticity equals 2, so medical expenses also increase by 30 percent. But now the increased

TABLE 2-3 Uncertain Medical Spending: Needs-driven or Innovation-driven?

<i>Scenarios (Responses to unexpected changes in percent)</i>			
	<i>Full coverage for retirees (1)</i>	<i>Capped benefits for retirees (2)</i>	<i>Efficient risk sharing (3)</i>
Panel 3A: Inelastic increase in medical needs			
Consumption spending in retirement	6.6	0.0	1.8
Consumption spending in working age	-1.6	-1.4	-1.4
Nonmedical Cons. in retirement	0.0	-8.5	-6.2
Nonmedical Cons. in working age	-6.4	-6.1	-6.2
Disposable income in working age	-3.6	0.0	-1.0
Capital investment	-8.0	2.9	0.0
Medical cost coverage	100	0	27
Panel 3B: Elastic response to medical innovations			
Consumption spending in retirement	6.6	0.0	-0.7
Consumption spending in working age	0.0	0.5	0.6
Nonmedical Cons. in retirement	0.0	-6.6	-7.4
Nonmedical Cons. in working age	-3.9	-3.4	-3.4
Disposable income in working age	-3.6	0.0	0.4
Capital investment	-11.5	-1.2	0.0
Medical cost coverage	100	0	-11

Source. Author's calculations.

Notes. Panel A assumes medical expenses are inelastic/needs-driven and increase 30 percent. Panel B assumes a 30 percent decline in quality-adjusted health care prices and an elastic response of medical consumption. Responses are computed for the calibrated model with stylized US policy, as detailed in the text. Col. 1: Retiree medical expenses are fully government funded on the margin. Col. 2: Retiree medical benefits are held constant. Col. 3: Efficient risk-sharing.

spending reflects a relative price change favorable to retirees. Full Coverage (Column 1) raises retiree consumption, reduces working-age consumption and capital investment as savers expect full health coverage in retirement. Capped Benefits (Column 2) leaves retiree spending unchanged and triggers an individually optimal shift from nonmedical to health care consumption. Workers also shift away from nonmedical to health care consumption, and they reduce savings slightly because lower-cost health care yields greater benefits in retirement than during working age.

Efficient risk-sharing (Column 3) now calls for retirees to share the positive effect of lower health care prices with working-age cohorts by accepting reduced transfers. Retirement spending declines, nonmedical consumption declines even more, whereas working-age income and consumption spending rise. The numerical values are again example specific, but the efficiency of a benefit reduction and the zero investment response are robust. These

results apply to any permanent price decline provided permanent retirees consume relatively more health care than younger cohorts.

International Risks: Wars and Foreign Assets

Wars. The risk of war is a major risk omitted in the analysis earlier. This is because burden-sharing is the standard paradigm of war financing, so risk-sharing does not provide much new insight. War expenses and war-related damages are negative shocks to the national resource constraint. A draft could be interpreted as a negative but temporary shock to working-age productivity. In an overlapping generations setting, efficiency would call for risk-pooling across generations. The results of a formal analysis would presumably resemble Barro's (1979) tax smoothing model.

However, the OG model gives different answers than that provided by the representative agent literature solving Ramsey tax problems. In representative agent models, the risk of war is best shared instantaneously through state-contingent claims (Lucas and Stokey 1983). In the OG model, efficient risk-sharing includes future generations. Their contributions to war finance cannot be collected instantaneously and requires taxation over time. State-contingent debt therefore does not avoid the need for a dynamic fiscal response to wars and other temporary shocks.

Foreign Assets. Foreign assets are the focus of international risk-sharing, as highlighted by Shiller (1993, 1999). He views international and intergenerational risk-sharing as alternative insurance mechanisms, and in his book (1993) he explains in some detail how international financial markets would help share income risks. In the same spirit, one could imagine macromarkets for country-specific fertility risks, longevity risks, and health care risks.

Unfortunately, the macromarkets that Shiller envisions do not yet exist. Today's financial markets facilitate capital flows and permit the trading of claims against capital income. For a quantitative perspective, note that the US States capital stock is about 225 percent of GDP, net foreign liabilities are about 25 percent of GDP, direct investments and equity holdings abroad are 36 percent of GDP, and foreign direct and equity investment into the USA total 43 percent of GDP (2003). Domestic residents thus retain a 190 GDP-percent exposure to domestic asset valuation risk ($=226-36$), almost equal to their 200 GDP-percent net wealth ($=225-25$). The low ratio of foreign assets to net worth ($43/200$) suggests that international assets are a smaller source of risk than any of the aggregate risks discussed earlier. While uncertainty about the return on foreign assets is a source of risk that might be shared across generations, it is unclear why international risk-sharing is so incomplete.

Endogenous Growth

A promising area for future research is intergenerational risk-sharing in an endogenous growth context, though most of the literature to date assumes exogenous growth. It has been noted that intergenerational risk-sharing could have much greater welfare benefits if it also affects economic growth. For instance Obstfeld (1994) posits that individuals can choose between low-risk, low-return and higher-risk, higher-return technologies. Under autarchy, risk-averse individuals choose relatively low-return investments. When risk are shared and thus reduced, individuals choose higher-return investments, leading to an increase in the *average* rate of growth. Applied in a multigenerational setting, the same argument would suggest that better intergenerational risk-sharing may encourage savers to select higher-risk, higher-return technologies, which would lead to an increase in average growth rates. The Obstfeld model suggests that growth effects could magnify the gains from risk-sharing by an order of magnitude.

Conclusions

This chapter has examined the allocation of major aggregate risks from an intergenerational perspective. Such risks cannot be eliminated by insurance, but they can be managed through risk-sharing. Government has a key role in this endeavor, because it can oblige future generations to participate in risk-sharing arrangements. Many fiscal institutions and practices can be interpreted as beneficial risk-sharing arrangements in this sense, but intergenerational risk-sharing is far from perfect. More specifically, we have reviewed risks related to productivity growth and to asset prices, demographic risks due to changes in fertility and longevity, and medical-expense risks created by changing health care needs and by innovations in medical technology.

With regard to macroeconomic risks, the message is to focus on productivity growth and not be distracted by uncertain asset values. Though uncertain asset values are indeed a risk for retirement savers and other asset holders, productivity growth is a much greater source of long run uncertainty. Fiscal institutions in the USA are well suited to provide safe claims to retirees via social security and government bonds, and asset valuation risk is shared via income taxes. Growth uncertainty has received comparatively little attention. Compared to a benchmark allocation with perfect risk-pooling, the retiree generation appears underexposed to productivity risk, though still overexposed to asset valuation risk. With regard to baby-boom and baby-bust phenomena, the key insight is that DB pensions and other intergenerational transfers have an important risk-sharing function. In an economy without such transfers, large cohorts are worse off than small cohorts. A large cohort's labor supply tends to depress wages and its

supply of retirement savings tends to depress asset returns. DB pensions impose relatively lighter burdens on larger cohorts and thereby help share demographic risk.

Fiscal policy has a similar risk-sharing role with regard to longevity. Longevity insurance is awkward from a distributional perspective, however. To share the financial risk of longevity, cohorts suffering from low-life expectancy have to make transfers to cohorts that enjoy a longer life. It is important therefore to condition risk-share on a trend path of rising longevity (e.g. by linking normal retirement to the longevity trend). If one focuses on deviations from the trend, one finds that annuitized pensions with defined benefits help share longevity risk across generations.

Uncertain health care expenses create similarly awkward insurance problems as longevity. Full insurance against medical innovations (say, though Medicare or Medicaid) mean that cohorts receiving relatively inferior care would have to transfer resources to cohorts that benefit from medical innovations. Such transfers would be implied, for example, by the fiscal balance rule of generational accounting. Efficient risk-sharing does not support full coverage for retirees' unexpected medical needs. We address this issue from two perspectives. One view is that medical needs are inelastic and separable from other consumption. In this case, efficient risk-sharing calls for the retiree generation's unexpected medical needs to be financed, in part, by taxes on future generations and, in part, by reduced nonmedical consumption. A different view is that medical care may be seen as a substitutable component of normal consumption; in this case, increased expenditures would be seen as driven by declining quality-adjusted prices and more than unit-elastic demand. In this case, efficient risk-sharing calls for the retiree generation to self-finance the impact of unexpected medical innovations, even for intergenerational transfers to decline a little. These results are, to reemphasize, macroeconomic and, therefore, disregard potentially challenging issues of cross-sectional distribution. They nonetheless provide a macroeconomic starting point for the design of health care policies.

Endnotes

1. Bernheim and Bagwell (1988) suggest an even stronger result: with intermarriage between dynasties, not even idiosyncratic risk would be an issue because altruistically motivated transfers from more fortunate to less fortunate family members would eliminate idiosyncratic risk. Barro's and Bernheim and Bagwell's papers specifically examine idiosyncratic taxes, but their insights apply equally to family responses to other idiosyncratic shocks.
2. The model abstracts from overlap between work status and income types. Much of the capital income during working age accrues in retirement funds or as housing wealth that is rarely liquidated before retirement. The economics of retirement savings is therefore not much distorted if one pretends that the return on working-age savings accrues at the time of retirement.

34 Henning Bohn

3. A temporary productivity shock would require a period of high growth to be followed by an offsetting period of low growth—a somewhat implausible scenario. Even then, consumption smoothing is insufficient under plausible assumptions to correct the overexposure of working-age cohorts to productivity risk (see Bohn (2004) for details).
4. Readers with a finance background may question this argument because up-front incentives imply a proportionally higher account balance at the individual level. If invested entirely in equities, this implies the same risk-exposure as an investment in an account without up-front incentives but tax-exempt earnings. The equivalence is invalid at the aggregate level, however, because up-front exemptions must be financed. *Ceteris paribus*, they imply a greater share of government debt in capital markets and hence in the average retiree's (enlarged) portfolio. On aggregate, savers must hold the capital stock and therefore cannot hold more equities in a system with up-front incentives than in a system with after-tax saving.
5. This uses the preferred error-corrections specification that exploits the stationarity of the dividend yield and the dividends/GDP ratio.
6. It is an open question to what extent estimates for the S&P500 generalize to other financial assets used for retirement savings. One may suspect that diversification into other asset classes would yield a lower valuation risk. The results for the S&P500 provide at least a starting point for thinking about long-run risks.
7. The main macroeconomic assumptions are a 30 percent capital share in production, an elasticity of intertemporal substitution of 0.5, a 25 percent share of old capital in the total return on retirement savings, an inelastic labor supply, a trend path with 1 percent annual population growth, 1.5 percent annual productivity growth, a 4.5 percent real return on capital, and a length of the retirement equal to one-third of work-life (15:45 years).
8. To be specific, bequests are modeled as 'accidental' due to stochastic mortality and imperfect annuitization. The latter is discussed below in the section on longevity risk.
9. See Bohn (1990), Shiller (1993), and Borensztein et al. (2005) for more discussion of GDP indexing.
10. Because stochastic inflation would create new risks, the point of this caveat is more to distinguish wage- and GDP-indexing from inflation-indexing than argue for nominal debt.
11. Note that variable immigration would have the same impact on cohort size as a variable fertility. The macroeconomic implications would be identical. My interpretation focuses on fertility mainly to avoid questions about whose preferences count, for the welfare analysis.
12. The factor two is conservative, to avoid overstating the differences. The ratio was 2.67 in 1999 for public and private health care expenses and about 5.7 for government-funded care. The overall value is relevant for welfare calculations. These ratios and the 65 percent coverage value below are computed from US Department of Health and Human Services, Centers for Medicare and Medicaid Services, National Health Care Expenditure Tables, online at <http://www.cms.hhs.gov/statistics/nhe/historical> downloaded 3/17/2005.

13. This percentage is about halfway between the Social Security Administration's intermediate cost and high-cost estimates for Hospital Insurance in 2035. It may be interpreted as standard deviation if one views the high- and low-cost estimates as $2\text{-}\sigma$ confidence bands.

References

- Altonji, Joseph, Fumio Hayashi, and Laurence Kotlikoff (1996). 'Risk Sharing Between and Within Families', *Econometrica*, 64(2): 261–94.
- Auerbach, Alan J., Lawrence J. Kotlikoff, and Willi Leibfritz (eds.) (1999). *Generational Accounting around the World*. Chicago, IL: University of Chicago Press.
- Barro, Robert (1974). 'Are Government Bonds Net Wealth?' *Journal of Political Economy*, 82(6): 1095–117.
- (1979). 'On the Determination of Public Debt', *Journal of Political Economy*, 87(5): 940–71.
- Bernheim, Douglas and Kyle Bagwell (1988). 'Is Everything Neutral?' *Journal of Political Economy*, 96(2): 308–38.
- Bohn, Henning (1990). 'Tax-Smoothing with Financial Instruments', *American Economic Review*, 80(5): 1217–30.
- (1999). 'Should the Social Security Trust Fund Hold Equities? An Intergenerational Welfare Analysis', *Review of Economic Dynamics*, 2(3): 666–97.
- (2001). 'Social Security and Demographic Uncertainty: The Risk-Sharing Properties of Alternative Policies', in John Campbell and Martin Feldstein (eds.), *Risk Aspects of Investment Based Social Security Reform*. Chicago, IL: University of Chicago Press, pp. 203–41.
- (2004). 'Intergenerational Risk Sharing and Fiscal Policy', Unpublished paper, University of California at Santa Barbara.
- Borensztein, Eduardo, Marcos Chamon, Olivier Jeanne, Paolo Mauro, and Jeromin Zettelmayer (2005). 'Sovereign Debt Structure of Crisis Prevention', IMF Occasional paper #237.
- Diamond, Peter (1965). 'National Debt in a Neoclassical Growth Model', *American Economic Review*, 55: 1126–50.
- Gokhale, Jagadeesh and Kent Smetters (2003). *Fiscal and Generational Imbalances*. Washington, DC: American Enterprise Institute Press.
- Kotlikoff, Laurence and Scott Burns (2004). *The Coming Generational Storm*. Cambridge, MA: MIT Press.
- Lucas, Robert and Nancy Stokey (1983). 'Optimal Fiscal and Monetary Policy in an Economy without Capital', *Journal of Monetary Economics*, 12: 55–93.
- Obstfeld, Maurice (1994). 'Risk-Taking, Global Diversification and Growth', *American Economic Review*, 84(5): 1310–29.
- Shiller, Robert (1993). *Macro Markets*. Oxford: Oxford University Press.
- (1999). 'Social Security and Institutions for Intergenerational, Intragenerational and International Risk Sharing', *Carnegie-Rochester Conference Series on Public Policy*, 50: 165–204.
- Smetters, Kent (2001). 'The Effect of Pay-When-Needed Benefit Guarantees on the Impact of Social Security Privatization', in John Campbell and Martin Feldstein (eds.), *Risk Aspects of Investment Based Social Security Reform*. Chicago, IL: University of Chicago Press, pp. 91–105.

36 Henning Bohn

United States Department of Health and Human Services, Centers for Medicare and Medicaid Services (2005). *National Health Care Expenditure Tables*, <http://www.cms.hhs.gov/statistics/nhe/historical> downloaded 3/17/2005.

Welch, Finis (1979). 'Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust', *Journal of Political Economy*, 87(5): S65–97.