

Propagation and Smoothing of Shocks in Alternative Social Security Systems

Alan Auerbach
University of California, Berkeley

Lorenz Kueng
Northwestern University

Ronald Lee
University of California, Berkeley

Yury Yatsynovich
Rensselaer Polytechnic Institute

Revised, February 13, 2018

Abstract

Even with well-developed capital markets, there is no private market mechanism for trading between current and future generations. This generates a potential role for public old-age pension systems to spread economic and demographic shocks among different generations. This paper evaluates how different systems smooth and propagate shocks to productivity, fertility, mortality and migration in a realistic OLG model. We use reductions in the variance of wealth equivalents to measure performance, starting with the existing U.S. system as a unifying framework, in which we vary how much taxes and benefits adjust, and which we then compare to the existing German and Swedish systems. We find that system design and shock type are key factors. The German system and the benefit-adjustment-only U.S. system best smooth productivity shocks, which are by far the most important shocks. Overall, the German system performs best, while the Swedish system, which includes a buffer stock to relax annual budget constraints, performs rather poorly. Focusing on the U.S. system, reliance solely on tax adjustment fares best for mortality and migration shocks, while equal reliance on tax and benefit adjustments is best for fertility shocks.

Keywords: notional defined contribution systems, pay-as-you-go systems, generational incidence

JEL codes: H55, J11

This research was supported by the U.S. Social Security Administration through grant #10-M-98363-1-01 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. We thank two anonymous referees for their helpful comments on earlier drafts. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA or any agency of the federal government.

1. Introduction

While the main function of old-age pension systems is to provide resources to elderly retirees, these systems can satisfy many other important government functions as well. Indeed, in circumstances where access to capital markets is good and many individuals can, alone or in conjunction with private employers, save for retirement, broad-based public pension systems may not be needed simply to provide retirement income and their other functions may take on greater prominence. One such function is the allocation and spreading of economic and demographic shocks among generations. Even with well-developed capital markets and informal family arrangements, there is no private market mechanism for trading between current and future generations, leaving government policy as the only broad-based option. A range of government policies, including national debt management, infrastructure investment, and public education expenditures, have important intergenerational consequences, but the size and variety of public pension schemes make them a natural place to focus for intergenerational policy.

Like private defined-contribution pension arrangements, funded defined-contribution public pension schemes result in one particular allocation of economic and demographic shocks among generations. For example, a demographic shock that leads to one age cohort being large relative to others will lead that cohort to experience relatively low lifetime wages (because of its high labor supply) and relatively low rates of return on its retirement saving (because of its high demand for retirement assets). But public schemes may deviate from the defined-contribution approach with respect to two criteria: asset accumulation and determination of contributions and benefits. With respect to the first criterion, systems may adhere to some form of strict pay-as-you-go (PAYG) approach, or to a more flexible approach that allows a fluctuation in the system's financial assets or liabilities within some stable range. With respect to the second

criterion, systems may adjust either contributions or benefits to maintain financial stability, and when adjusting benefits may adjust them immediately or in the future.

This paper relates to a large literature that studies how aggregate and idiosyncratic risks are shared in models with inter-generational linkages, starting with Diamond (1977).¹ Our main contribution to this literature is to analyze how specific demographic and economic shocks policymakers are concerned about propagate through different versions of public pension systems currently employed or proposed in different countries.

Two earlier papers (Auerbach and Lee 2009, 2011) studied a variety of existing and hypothetical unfunded arrangements. These included the actual and hypothetical Swedish systems, the actual German system, and three stable variants of the existing U.S. social security system, evaluated according to a variety of welfare criteria, such as internal rates of return and an approximation of expected utility. Our findings, particularly in the second paper, suggested that the methods of spreading shocks across generations can have significant effects on welfare. But questions remain about the channels through which these effects operate.²

Understanding the effects of an existing or proposed system on welfare is, ultimately, our objective in studying the spreading of shocks. However, we would also like to understand why certain systems seem to perform better in the welfare dimension and how other potential systems would perform in response to different patterns of shocks. Our past welfare analysis was based on empirically estimated demographic and economic stochastic processes for the United States,

¹ This literature is too large to summarize adequately here and we refer the reader to the excellent surveys by Feldstein and Liebman (2002) and Attanasio et al. (2016). Two more recent extensions of the overlapping generations (OLG) framework developed by Diamond (1977) worth mentioning in our context are Shiller (1999) and Bohn (2009), who show that public pension systems can reduce consumption risk of all generations.

² Other studies that analyze specific designs of public pensions include Matsen and Thøgersen (2004), who investigate the optimal mix between PAYG and fully funded systems, Krueger and Kubler (2006), who study the welfare benefits of a minimum pension, and Ludwig and Reiter (2010), who ask how pension systems perform in smoothing fertility shocks.

but patterns in the future or in other countries may differ; it would be useful to have a more general picture of how different systems perform in response to different types and patterns of shocks. Using the stochastic modeling approach of our previous work, by looking only at the particular shocks of interest, one at a time, is difficult because each type of shock has complex economic effects and channels that cannot be determined without an explicit general equilibrium model. Thus, we utilize a modified version of the Auerbach-Kotlikoff (1987) dynamic general-equilibrium OLG model that incorporates realistic patterns of fertility and mortality and shocks to productivity, fertility, mortality and migration. While this model is well adapted for our purposes, it is worth highlighting some limitations of our analysis.

One issue concerns the fact that, although we model the responses of the economy to a range of shocks, the model is effectively deterministic, in the sense that households do not anticipate the shocks and respond to them after they occur. In particular, due to the complexity of our model economy and the particular counterfactuals we study, we abstract from modeling aggregate or idiosyncratic uncertainty. Although recent studies have made substantial progress in incorporating aggregate shocks into realistic large-scale OLG models, these models are still too limited to deal with our setting. For instance, Hasanhodzic and Kotlikoff (2013) incorporate aggregate productivity shocks in a large-scale OLG model with a Social Security system that is simpler than ours and does not incorporate demographic changes. Rios-Rull (2001) on the other hand incorporates stochastic population changes but does not include government and hence cannot study the role social security plays in propagating or smoothing of these shocks. The latter study, like much of the literature analyzing aggregate shocks such as Krueger and Kubler (2006) and Ludwig and Reiter (2010), also analyzes such shocks using a quadratic approximation of the full rational expectations equilibrium around a deterministic steady state.

Other studies incorporate idiosyncratic risks but not aggregate uncertainty. The models by Hong and Rios-Rull (2007) and Nishiyama and Smetters (2007) for example feature a simplified version of social security and idiosyncratic demographic shocks, but no aggregate uncertainty. The latter study finds that even allowing for idiosyncratic shocks “significantly increases the complexity of the model and the required computation time from several hours to typically several days per simulation.” While we do incorporate uncertainty with respect to mortality in our model, we couple it with an assumption of complete annuity markets, so that household saving decisions are not affected.

One limitation of our study is therefore the absence of precautionary savings as an element of the household’s intertemporal optimization, and we leave such an analysis to future research. However, it is worth pointing out that since our study focuses on the comparison of the dynamic responses across different social security systems, it is not obvious whether and to what extent the existence of precautionary savings would change these relative comparisons. A second limitation, associated with the absence of idiosyncratic uncertainty, is that while we can and do evaluate the spreading of shocks across generations, we cannot analyze risk-sharing within generations. Finally, in studying shocks around the steady state of an economy, we limit our consideration to shocks that, though potentially of very long duration, are not permanent.

Our analysis yields several interesting results. First, for our calibration based on U.S. historical experience, productivity shocks are the most important source of welfare volatility, so success at smoothing such shocks could be quite valuable. Yet no system is particularly effective at smoothing productivity shocks, because the tax cuts or benefit increases that positive productivity shocks make feasible generally aid cohorts already gaining from the productivity shock. However, among variants of the U.S. social security system, reliance on benefit

adjustment smooths such shocks most effectively, as retirees' direct gains from the shocks themselves are smaller than are those of workers.

Adjustments in response to fertility shocks, the second most important shocks, are more effective when relying on a mix of tax and benefit adjustments, because cohorts hurt by the initial shock in different ways (a decline in wages from an increased labor force vs. reduced consumption due to larger families) are helped differentially by tax and benefit adjustments. For mortality and migration shocks, which in our calibration are the least important and generate roughly equal amounts of volatility in well-being, tax adjustments are better at smoothing, because these shocks affect workers most – migration shocks because of increased labor-market competition, mortality shocks because of less need for resources among the elderly. Also, social security has a greater ability to smooth shocks in these cases than for productivity and fertility shocks, because the initial effects of the shocks themselves are less widespread.

An important lesson from our analysis is that, even where smoothing is possible, some combinations of tax and benefit adjustments fail to smooth shocks. In fact, some social security systems actually concentrate the effects of shocks for certain parameterizations and shocks. For example, while reliance primarily on tax adjustments effectively smooths the migration shock, this is not the case when the primary adjustment is through benefits.

Finally, considering the performance of the German and Swedish systems, we find that the German system performs well overall relative to the U.S. system, successfully combining tax and benefit adjustments in smoothing different shocks. The Swedish system, on the other hand, performs rather poorly except in the case of the fertility shock, even though it alone has a buffer stock mechanism that relaxes the requirement of short-run cash-flow balance.

2. The Model

The model we use is adapted from that laid out in Auerbach and Kotlikoff (1987, chapter 11) and used subsequently by Auerbach et al. (1989) to evaluate the economic effects of public pension systems in several countries. That original model was a perfect foresight, dynamic general equilibrium model with variations in fertility that permitted analysis of the interactions of demographic transitions and different public pension systems. However, several modifications are needed to make that model suitable for the current task. In particular, the model had a very simplistic approach to fertility, assuming that it was concentrated at one age, had no individual uncertainty as to life expectancy, no migration, and assumed a smooth rate of productivity growth. We will go through the newly adapted version of the model, indicating how we use it to measure the effects of shocks to productivity, fertility, mortality, and migration.

In the model, individuals live for up to T years, spending the first 20 as minor dependents of parents who make consumption decisions on their behalf. At any given time, the household consists of one parent and any minor children. Household utility in each year is based on the parent's consumption and leisure and each child's consumption. The household maximizes lifetime family utility subject to a lifetime budget constraint. Each age cohort of adults consists of individuals who are identical, *ex ante*, but who have different mortality experience.³

For simplicity, we assume that live children are born to parents between the ages of 21 and 40, and that mortality begins at age 60, after children have left the home. This means that there are no orphan children, which would add complexity to the model.⁴ We assume that births

³ Note that while we assume a distribution of births across childbearing ages, we treat these births as deterministic, with the representative household having this entire distribution.

⁴ Although relaxing this assumption would be desirable in principle, it would not be of major quantitative importance. In the US life table for 2005-2009, only 9 percent of female births die before age 60, and of those who survive to begin childbearing at age 20, only 8 percent die before 60. The probability that both parents would die before 60 in a more realistic model with two parents for each child would be less than 1 percent.

follow a baseline age-specific fertility profile z_a between ages 21 and 40, which may be shifted by an AR(1) shock ζ_t so that fertility at age a and date t – the number of children born to a household of age a in year t – is $f_{at} = z_a \zeta_t$. The baseline mortality hazard rate for age a at date t is m_a . We assume that this vector of age-specific mortality rates can be hit by a multiplicative shock that also follows an AR(1) process, μ_t , so that the survival probability hazard for age a and date t is $s_{at} = (1 - m_a \mu_t)$. We assume, for the sake of simplicity, that there are no trends or aggregate uncertainty in the probability age profiles of either fertility or mortality.

As to intergenerational linkages after children become adults, we assume that there are no inter vivos gifts or intentional bequests. Given uncertain lifetimes, though, individuals dying before age T could still leave accidental bequests. Rather than dealing with accidental bequests at different ages, which would make solution of the model extremely complicated, we assume perfect Yaari annuity markets, so that individuals fully annuitize their retirement savings and therefore leave no bequests regardless of when they die. This means that the adjusted return to saving should equal $(1 + r_t^n) / s_{at}$ for age- a individuals at date t , where r_t^n is the net, after-tax rate of return on capital at date t . Note that the combination of mortality and perfect annuity markets should leave the household optimization problem unaffected, as higher rates of return on annuities will just offset higher discount rates induced by mortality. That is, even though the household's objective function now incorporates expected mortality, we can determine the household's optimal planned consumption and labor supply paths (contingent on survival) ignoring mortality in both the objective function and the rate of return.

Figure 1 displays our baseline fertility and mortality hazard probabilities. These are for the United States, taken from the Human Mortality Database and the Human Fertility Database

for 2010,⁵ except that, to accommodate modeling assumptions, values for fertility are set to zero below age 21 and above age 40, and values of mortality are set to zero for ages below 60.⁶

We assume that the household maximizes a lifetime family utility function that is time-separable, separable across individuals, and having a nested CES structure for adults within periods (between consumption and leisure) and between periods. Taking fertility and mortality into account, the household's objective function at age 21 may be written:

$$(1) \quad \sum_{a=21}^{a=T} \frac{(1+\delta)^{-(a-21)}}{1-1/\gamma} \left(\prod_{i=60}^{i=a} s_{i,t+i-21} \right) \left\{ (C_{a,t+a-21}^{1-1/\rho} + \alpha l_{a,t+a-21}^{1-1/\rho})^{\frac{1-1/\gamma}{1-1/\rho}} + \sum_{k=\max(1,a-D+1)}^{k=\min(a-20,G)} f_{a-k+1,t-k+1} \omega_k H(a)_{k,t+a-21}^{1-1/\gamma} \right\}$$

where $C_{a,t}$ is adult consumption at age a in year t , $l_{a,t}$ is the corresponding leisure, $H(a)_{k,t}$ is the consumption of a child of age k in year t for a parent of age a , T is maximum life (set to 100 in our simulations), D is the maximum age of child-bearing (here, assumed to be age 40), and G is the age after which children are adults and leaders of their own families (here, assumed to be 20). As in Auerbach and Kotlikoff (1987), the terms ω_k are weights of children in the utility function which are assumed to increase linearly from 0.25 at age 1 to 0.50 at age 20, i.e., $\omega_k = .25 + .25*(k-1)/19$. In expression (1), there are also four household preference parameters: δ is the pure rate of time preference, α is the leisure intensity parameter, ρ is the intra-temporal elasticity of substitution between consumption and leisure, and γ is the inter-temporal elasticity of substitution over consumption (and, in the case of adults, leisure as well) at different dates.

Following the calibration in Auerbach et al. (1989), we set the last three parameters equal to 1.5, 0.8, and 0.35, respectively. We choose the final parameter, δ , to target the historical postwar

⁵ The data are available at <http://www.mortality.org/> and <http://www.humanfertility.org/cgi-bin/main.php>.

⁶ We inflate the remaining fertility profile to offset the excluded years of fertility in order to produce the same number of births per adult as for the full fertility profile. We do not adjust the mortality profile, which gives us the correct measure of life expectancy at age 60, although overstating life expectancy at earlier ages.

U.S. before-tax rate of return to capital measured in Auerbach, Kotlikoff and Koehler (2016), just under 6.5 percent. The resulting value for δ is -.015.⁷ The household maximizes utility as defined in (1) subject to a budget constraint that equates the present value of household consumption to the present value of after-tax labor income plus pension benefits.

The economy has one production sector, with a representative competitive firm producing output subject to a constant-returns-to-scale Cobb-Douglas function in capital and labor. Hence there are no pure profits, with returns to capital and labor exhausting the firm's income. The economy is closed in the initial simulations we consider, so the production sector's capital stock is determined by household plus government asset accumulation. Labor equals the sum of labor supplied by cohorts of different ages, where we assume that different ages of labor are perfect substitutes but differentially efficient as described by an empirically estimated age-based efficiency profile, e_a , also taken from Auerbach and Kotlikoff (1987).

We assume that individuals work starting at age 21, with the retirement age endogenously determined (as the age at which zero labor supply is optimal) by preferences and factor prices. This retirement age is distinct from the initial age of benefit receipt, which we set at 67 for all social security systems considered, consistent with the normal retirement age under the current U.S. system once it is fully phased in. There is a deterministic trend in total factor productivity growth, at a rate of 1.5 percent per year, which is implemented in two pieces: greater steepness in the efficiency profile, e_a , plus a trend in the labor endowment, which increases equally the individual's efficiency at supplying labor and in consuming leisure. This method produces the right wage profile for each cohort but also avoids any trend in the work/leisure ratio, as discussed in Auerbach et al. (1989). We assume the presence of multiplicative productivity shocks around

⁷ One may rationalize this low time preference rate as offsetting the model's lack of a precautionary saving motive (which would be present without a complete annuities market) and possibly also the lack of bequests.

a steady state value, with these shocks again following an AR(1) process. These productivity shocks will affect the market returns to labor and capital at each date t , w_t and r_t .⁸

While there is no immigration in our model as described so far, we can consider the effects of a one-time migration shock by introducing a sudden, one-time inflow of individuals with an age profile based on that of existing immigrant flows. We assume immigrants enter the social security system upon arrival, with no credit for earlier years, and that they bring no capital but have the same productivity as domestic workers, as is necessary in our model with identical individuals in each age cohort.⁹

We calibrate the various shocks as follows: for fertility, mortality, and productivity, we assume a one-time shock followed by decay based on an AR(1) process. Based on estimates using postwar U.S. data (details available on request), we set the AR coefficient equal to 0.9 for each of these processes, and consider the effects of the largest of each type of shock experienced during the period, which (in our model) translates into a productivity shock of 4 percent, a fertility shock of 18 percent, and a mortality shock of 4 percent, all relative to their respective baseline levels.¹⁰ Figure 2 shows the evolution over time of the population age structure after a shock that increases the fertility rate at every child-bearing age by 18 percent. Starting with a smooth initial population structure, there is a jump up in the young population by year 10. By

⁸ While productivity shocks hit both returns to capital and wage rates, it would also be straightforward to evaluate separate shocks to the two processes by introducing another shock that affects only wages, via a shift in the age-wage productivity profile, e_a , or only returns to capital, via stochastic depreciation. We have considered the effects of adding stochastic depreciation to the model, following the calibration of Ambler and Paquet (1994) in terms of the relative size of productivity and depreciation shocks and the latter being serially uncorrelated. The spreading of such shocks is nearly the same under all variants of social security we analyze, and hence has little effect on our conclusions. Therefore, we do not present these results below. We conjecture that this finding relates to the depreciation shock being short-lived but leave additional consideration to future research.

⁹ Having different economic state variables than domestic residents of the same age, immigrants choose different consumption and labor supply paths, even though we assume they have the same labor productivity and preferences.

¹⁰ Because the model is nonlinear, one cannot simply scale the results down proportionally to estimate the effects of smaller shocks, but we found little qualitative impact on our results for shocks half as large.

year 50, there is a broader increase in the population showing both the initial fertility shock and a smoothed echo effect of this initial shock. In the final steady state, the population, shown in Figure 3, is larger as a consequence of the shock, but the original age structure is restored.

Figure 3 also shows the comparable population evolution in response to a shock that increases mortality rates uniformly across all ages by 4 percent. This shock reduces the elderly and hence overall population by year 10, but the reduction disappears over time as mortality rates return to their original levels. Unlike in the case of a shock to fertility, where the population is permanently higher, under a mortality shock (given our assumption that changes in mortality occur only after age 60 and therefore do not interact with child bearing) the population has returned in the final steady state to its pre-shock pattern and level. For migration, we consider a large shock based on recent history. In particular, we assume a positive immigration shock equal to 0.5 percent of population, roughly the order of magnitude of the increase in net immigration into Germany in 2015 over recent historical levels.¹¹ Figure 4 shows the profile we use, based on immigrants to Germany in 2014.¹² The impact on total population, relative to the no-shock trajectory, appears in Figure 3, along with those of the other two demographic shocks.¹³

Finally, we include the government sector in the model, consisting of two components, general government and the public pension system. General government follows a parsimonious specification of government purchases as consisting of age-based and non-age-based components, with the age-based components (e.g., education spending, old-age medical care) held constant (except for trend productivity growth) relative to their respective population groups and non-age-based components (e.g., defense spending) held constant (except for productivity

¹¹ <https://www.destatis.de/EN/FactsFigures/SocietyState/Population/Migration/Tables/MigrationTotal.html>.

¹² These data come from Eurostat, http://ec.europa.eu/eurostat/en/web/products-datasets/-/MIGR_IMM2CTZ.

¹³ Note that, unlike in the case of a fertility shock, the population level overshoots in the case of the migration shock because a greater share of the immigrant population is of childbearing age.

growth) relative to total population. We break spending down into age-specific and non-age-specific categories and hold spending for each category, i , constant at g_i per member of the relevant population, N_i , for $i = y, m, \text{ and } o$ (young, middle-aged, and old) or for the total population. That is, overall general government spending at any date t equals:

$$(2) \quad GOV_t = g_y N_{yt} + g_m N_{mt} + g_o N_{ot} + g_n (N_{yt} + N_{mt} + N_{ot})$$

We solve for the values of g_i by entering relative values of spending (gN) in each of the four categories, taken from Auerbach and Kotlikoff (1987, chapter 11) – .306, .172, .141, and .381, respectively – and then scaling them so that government spending equals the exogenously determined level of revenue in the initial steady state. During the transition, we keep these government spending weights, g_i , constant except for trend productivity growth, so that government spending grows smoothly except for changes due to demographic shifts.

General government is funded by proportional income and consumption taxes; we ignore the use of government debt for the general government budget. In the initial steady state, we set the income tax to 20 percent and the consumption tax equal to 3 percent, which are similar to recent estimates for the U.S. tax system from Auerbach (2002). Because shocks may cause required spending and revenue for given tax rates to fluctuate, we let the consumption tax rate vary to ensure annual budget balance for the general government.¹⁴

a. Modeling Public Pensions

In our analysis, we begin with a consideration of the actual U.S. social security system, in which we assume that there is year-by-year cash-flow balance but in which there may be

¹⁴ For the productivity shock, this tax rate falls to 1.5 percent (due to a higher tax base) before rebounding. For the fertility shock, the tax rate rises to 3.9 percent (to cover higher government spending on the young) before declining. For the smaller mortality and fertility shocks, the tax rate fluctuates narrowly between 2.97 and 3.02 percent.

different assumptions with respect to the mix of benefit and tax adjustments used to maintain annual cash-flow balance in the presence of shocks.¹⁵

For any version of the U.S. system, we calculate benefits roughly as under current law, taking the average of past labor earnings, inflated by wage growth between the date of earnings and the date of benefits receipt to calculate average indexed monthly earnings.¹⁶ We then solve for the payroll tax needed for budget balance, according to the expression,

$$(3) \quad \theta_t \sum_{j=21}^T w_{j,t} (1 - l_{j,t}) N_{j,t} = R_t \sum_{j=67}^T AIME_{t-j+1} N_{j,t}$$

where θ_t is the payroll tax in year t , $w_{j,t}$ is the wage rate of an age- j individual in year t , $N_{j,t}$ is the population of age j in year t , $AIME_t$ is the average indexed monthly earnings for an individual born in year t , and R_t is the replacement rate in year t . We set $\theta = .106$ in the initial steady state, consistent with the current OASI portion of the payroll tax for the United States, and solve for R according the expression (3). During the transition, when shocks occur, we adjust either R , θ , or both annually to ensure that (3) continues to hold, with the relative adjustments of R and θ depending on the fraction of the budget adjustment assumed to come through benefits.

After consideration of variants of the U.S. system, we will also evaluate systems in Germany and Sweden. For the German system, which also basically requires annual cash-flow balance, we follow the description in Auerbach and Lee (2011). Each beneficiary at date t receives the same benefit, B_t , so that budget balance requires that

¹⁵ This is a simplification of the actual working of the U.S. social security system, which maintains a small trust fund and by law requires no adjustment until the trust fund is exhausted, at which time benefits must fall to maintain a zero trust fund balance. Our modeling of different tax and benefit adjustments is meant to reflect alternative outcomes of the political process rather than simply the rules of the existing system.

¹⁶ For simplicity, we include all years of work in this calculation, rather than the 35 years with highest earnings, as currently used for the US system.

$$(4) \quad \theta_t \sum_{j=21}^T w_{j,t} (1 - l_{j,t}) N_{j,t} = B_t \sum_{j=67}^T N_{j,t}$$

The benefit itself evolves from one year to the next according to the following formula:

$$(5) \quad B_t = B_{t-1} \cdot \frac{w_{t-1}(1-\theta_{t-1})}{w_{t-2}(1-\theta_{t-2})} \cdot \left[1 - .25 \cdot \left(\frac{PR_{t-1} - PR_{t-2}}{PR_{t-2}} \right) \right]$$

where w_t is the aggregate wage in year t , normalized for age-specific productivity differences, and PR_t is the pensioners' ratio in year t , which we define to be the ratio of the adult population above the age of pension receipt (age > 66) to those of working age (ages 21-66). While expression (5) determines how benefits evolve over time, it does not fix the level of benefits. To facilitate comparison with the U.S. systems, we fix B in the initial steady state so that expression (4) is satisfied by the same payroll tax rate as is assumed for the U.S. systems. During the transition, the benefit evolves according to expression (5) and the tax rate θ is determined so that expression (4) is satisfied. Thus, both B and θ will change from year to year during the transition, and so the German system resembles a variant of the U.S. system that adjusts both taxes and benefits, although not necessarily one for which the adjustment shares are constant.

The other system we consider, a variant of the existing Swedish system, differs from the others in that (1) only benefits are adjusted; and (2) there is a trust fund that serves as a buffer stock. It is a notional defined contribution (NDC) system, in which each worker accumulates *notional pension wealth* (NPW) in a virtual account and then receives an annuity based on NPW accumulated as of retirement. The system uses a rate of return that is intended to be more sustainable for an unfunded system, equal to the growth rate of the wage level, and has an additional mechanism, referred to as a balancing mechanism, to adjust this rate of return further if a measure of system sustainability based on explicit and implicit assets and liabilities deviates

from its target ratio. The balancing mechanism in the actual Swedish system is asymmetric, lowering the rate of return when the funding ratio is low but never increasing the rate of return when the funding ratio is high, but we modify this in our analysis in order to generate a stable steady state for the system and to allow comparison with the other systems, which have no trust fund accumulation. We also set the rate of return equal to the growth rate of wages rather than of the wage level, to take account of labor force growth. A sketch of our modified version of the Swedish system follows. Auerbach and Lee (2011) provide further details.

During working years, individual i accumulates notional pension wealth at the beginning of period t , NPW_{it} , according to the formula:

$$(6) \quad NPW_{it} = NPW_{it-1} \cdot (1 + r_{it-1}) + T_{it-1}$$

where T_{it-1} equals the individual's payroll taxes at the end of period $t-1$ and r_{it-1} is the system's notional rate of return in period $t-1$. In the "normal" regime (when the balancing mechanism doesn't bind), r equals the rate of growth of total wages during the same period, g .

Upon reaching the retirement age, R , which we will mark by the exogenous age of benefit receipt used in all of our systems, the individual gets an annuity based on the value of that individual's notional pension wealth at that date; if i is indexed by the individual's year of birth, this would be $NPW_{i,i+R}$. The annuity payment at the end of year $i+R$ equals:

$$(7) \quad x_{i,i+R} = NPW_{i,i+R} \left[\sum_{t=i+R}^{i+T} (1 + \bar{g})^{-(t-i-R+1)} P_{i+R,t} \right]$$

where \bar{g} is the long-run average growth rate of wages, g , $P_{i+R,t}$ is the survival probability of generation i between the retirement date and date t , and T is the maximum lifespan. Since realized wage growth in any given year will not equal \bar{g} , the annuity payment at the end of year $i+R+1$ is adjusted up or down by the difference between \bar{g} and realized growth, g_{i+R+1} , between

year $i+R$ and year $i+R+1$. Thus, the actual annuity will be level if the growth rate actually equals \bar{g} , but otherwise it will grow or fall each year according to the realized value of g_t .

Because there is nothing built into the benefit calculation just described that ensures that benefits and taxes will be equal in any given year, the system can accumulate debt or assets. The actual system has a balancing mechanism that is activated when a “balance ratio”, b_t , falls below a threshold of 1, while we also apply the mechanism whenever the balance ratio is above 1.

The balance ratio (suppressing subscripts, since all are in the same year) is defined by:

$$(8) \quad b = \frac{F+C}{NPW+P}$$

where F equals financial assets accumulated in the system, positive or negative; NPW is aggregate notional pension wealth as of that date for all individuals below age R ; P is the present value of remaining annuity payments to all retirees based on the discount rate \bar{g} (that is, assuming that each retiree’s annuity remains constant at its current value, and discounting these level payments using the gross discount factor $(1 + \bar{g})$); and C is the so-called “contribution asset,” meant to approximate the present value of future tax payments by non-retired current participants. It equals the three-year moving average value of tax payments times the three-year moving average of “turnover duration,” which equals the average length of time between tax payments and benefit receipt. We calculate the turnover duration in year t as the average (weighted by benefits) age of benefit receipt in year t minus the average (weighted by taxes) age of tax payment in year t . Note that four measures, F , C , NPW , and P , are all backward-looking.

When the balancing mechanism is in effect (in our modeling, whenever $b \neq 1$ and so essentially always outside the steady state), the credited gross rate of return is not $(1+g_t)$ but

$(1+g_t) b_t$ and the annuity in each year is adjusted up or down from the previous year by $[(1+g_t) b_t - (1+\bar{g})]/(1+\bar{g})$ rather than $(g_t - \bar{g})/(1+\bar{g})$.

In a steady state, the Swedish system must have constant values of the balance ratio, b , and the ratio of financial assets to some measure of aggregate activity, e.g., total wages, F . As we have modeled the system, a steady state occurs when $(F, b) = (0, 1)$. That is, the system is such that, with a constant growth rate and no trust-fund assets ($F = 0$), taxes and benefits will be equal and the balance ratio, b , will equal 1. The results presented below are for parameterizations where this “normal” steady state in the Swedish system is unique and stable. But there is no assurance of stability or uniqueness more generally.¹⁷

For all public pension systems, we assume that individuals perceive some linkage between social security benefits and contributions, that is, that a portion of payroll taxes are viewed as being offset by the incremental benefits they generate. The higher the perceived tax-benefit linkage, the lower the labor supply distortion caused by the payroll taxes. We follow Auerbach and Kotlikoff (1987) and set the tax-benefit linkage at 0.25, meaning that individuals ignore one fourth of payroll taxes when making labor supply decisions.¹⁸

¹⁷ Indeed, for many reasonable parameter combinations, this equilibrium is unstable, and the only stable steady state is one for which $F > 0$ and $b > 1$. For our base case parameter assumptions, for example, stability is just satisfied – a small increase in the pure rate of discount, δ , leads to instability. Stability analysis of a simpler version of the model (with three periods and no feedback effects of the capital-labor ratio on factor prices) confirms based on a computation of eigenvalues that the “normal” equilibrium $(F, b) = (0, 1)$ is stable only for interest rates below some critical threshold. For higher interest rates, a perturbation starting at this equilibrium pushes the economy toward the alternative equilibrium, for which trust fund assets may be quite high and, as a result, the taxes needed to finance benefits quite low. A possible intuition for this result is that, when a shock pushes taxes and hence trust fund assets above zero, these assets then accumulate faster than the increase in benefits, with this process continuing until benefits get sufficiently high to absorb the increasing trust fund interest. This suggests that there may be powerful forces in the direction of asset accumulation within the actual Swedish system, where the balancing mechanism does not apply for shocks that increase trust-fund assets and drive the balance ratio above 1, although this is only a conjecture given the various simplifying assumptions made in the construction of our model of the economy.

¹⁸ While one might expect there to be differences in linkage among the systems, based on their transparency and within-cohort redistribution, we lack evidence on such variation and do not anticipate that such differences would exert a major influence in the extent of smoothing of shocks.

b. Solution of the Model

To consider the solution of the model, first assume that there are no shocks to mortality, fertility or productivity. In this case, the economy eventually follows a steady state path, so we start by solving for this steady state, using the Gauss-Seidel solution technique laid out in Auerbach and Kotlikoff (1987). Now, suppose that the economy is initially in this steady state, in year 1, and is then hit by one of the four types of shocks in year 2, with no further shocks thereafter (but the fertility, mortality and productivity shocks fading out only gradually in accordance with the AR(1) specifications for each type of shock). Since the shocks eventually die out completely, the economy will gradually return to the same steady state, assuming that the shock does not induce any permanent change in the social security system.¹⁹

To solve for the transition path, we assume that the shocks occur by surprise in year 2, after which all agents in the economy are endowed with perfect foresight. Thus, transition back to the steady state corresponds to a perfect foresight transition path, along which the actual paths of all variables are taken into account in household and firm optimization decisions at each date.

c. Measuring the Smoothing of Shocks by Public Pensions

Once the transition path is determined, we can calculate the impact of each particular type of shock on each cohort's welfare. Based on the expected utility given in expression (1), we calculate the *wealth equivalent* of the cohort's utility change as the scalar that, when multiplied by the household's steady-state vector of consumption and leisure, equalizes steady state utility and actual utility along the economy's transition path in the presence of the shock. The wealth

¹⁹ Such changes do not occur in the U.S. system or in our variant of the Swedish system (because the balancing mechanism forces the balance ratio back to 1 and financial assets to zero). Under the German system, by contrast, there is no terminal condition to guarantee such invariance in the long-run system, since all that is determined is how benefits and taxes adjust in each year to maintain cash-flow budget balance. In our simulations, however, the post-transition tax rate in the German system is very close to the original one.

equivalent is an index that measures the money-metric loss or gain induced by the corresponding shock, scaled by the cohort's present value of resources in the absence of the shock. By comparing wealth equivalents for a given shock in the absence of social security and under a particular social security system, we gain an estimate of the extent to which that scheme mitigates or exacerbates the shock for a given cohort.

For shocks to productivity growth, the wealth equivalent as just described will give us a clear measure of changes in well-being due to the shocks. For shocks to fertility and mortality, however, the issue is more complicated, because the expected number of individual-years of consumption and leisure change. For example, higher fertility will increase the number of children-years and lead the household to shift more resources to children; higher mortality will reduce the resources that the household wishes to devote to consumption and leisure in later periods of adult life. These changes complicate comparisons of household well-being. For example, simply having more children would lower the measured level of utility (since $\gamma < 1$), even if every element of the vectors of consumption and leisure were the same.

To deal with this issue, we measure wealth equivalents using the demographic parameters that would apply in the absence of shocks. That is, we measure the utility of consumption and leisure profiles in transition relative to those in the absence of shocks using the mortality and fertility rates that would apply in the absence of shocks.²⁰ With this approach, the wealth equivalent for a household will exceed 1 if and only if the household's observed vector of consumption and leisure would be preferred to the bundle chosen in the absence of shocks, for the fertility and mortality profiles that would apply in the absence of shocks.²¹

²⁰ One could also perform the comparison holding the demographic parameters at their values in the presence of shocks. The key is to hold the parameters constant in the comparison.

²¹ Our approach to measuring relative well-being makes sense if changes in fertility and mortality are due to exogenous shocks, but the issue would be more complicated if fertility and mortality shocks reflect endogenous

To measure the effectiveness of any social security system at smoothing shocks, we use the variance of wealth equivalents, where we scale each cohort's wealth equivalent by the trend level of the economy (based on population and productivity growth), discounted using the rate of return to capital. The logic of using a reduction in variance to measure smoothing is that variance provides a second-order approximation of the welfare loss from fluctuations in outcomes for a concave social welfare function. Our scaling is intended to avoid having our measure of volatility depend on a simple shift in resources to earlier or later generations that a social security system might produce.²² Likewise, we do not focus on how different systems affect mean levels of wealth equivalents under different shocks, which would pick up changes in economic efficiency that are not our focus here.

3. Basic Results

The effectiveness of any social security system in smoothing shocks will depend on both the nature of the shocks and the system's response mechanism. It is well understood that shocks of longer duration will be less susceptible to smoothing, to the extent that a broad range of cohorts already experiences similar effects. But a related point is that, for a given duration, some shocks, for example a productivity shock, will affect a wider range of cohorts than, say, a mortality shock, since the former directly affects returns to labor and capital and all those who receive them, while the latter primarily affects only older cohorts. Finally, of course, social

behavior by the household. For example, fertility might change because of a reduction in the cost of raising children. To evaluate adequately the welfare effects of such a change, it would be necessary to incorporate household fertility decisions explicitly in our model.

²² That is, a transfer of resources from each person in one birth year yields $(1+r)/(1+n)(1+g)$ dollars per individual scaled by the initial level of individual resources in the next birth year, where r is the rate of return, n is the population growth rate, and g is the productivity growth rate. Thus, for a system that shifted resources across cohorts so that positive and negative shocks were shifted into the future, the simple variance of wealth equivalents would rise if $(1+r)/(1+n)(1+g) > 1$. This normalization has relatively little impact except for the Swedish system, which through its buffer-stock mechanism produces considerable shifting across cohorts for some shocks.

security systems differ in the extent to which they rely on taxes versus benefits in maintaining cash-flow balance, and this affects their ability to smooth shocks that hit different cohorts, as do differences in the way that benefit and tax adjustments occur. Shocks that hit younger cohorts might more easily be spread using taxes, just as benefits might be used to address older cohorts; and benefit adjustments, especially, differ in the way they hit different cohorts.

Figure 5 provides a summary evaluation of the different social security systems in terms of effectiveness at smoothing the various shocks considered. For each shock, the figure plots the variance of wealth equivalents. Because of differences in variances across shocks, we plot them using a logarithmic scale and provide further detail on the two largest shocks (to productivity and fertility) using different expanded scales in the lower panel of the figure. To illustrate the relative effectiveness of benefit and tax adjustments, we show results for versions of the U.S. system ranging from no benefit adjustment to no tax adjustment. The entries on the right in the figure are for the German and Swedish systems and for the benchmark case of no social security.

The figure includes a number of interesting results. First, the shocks vary substantially in their underlying importance, with the effects of productivity shocks an order of magnitude larger than those of fertility shocks and three orders of magnitude larger than those of mortality and migration shocks. Second, focusing initially on the different variants of the U.S. system, we see that the relative effectiveness of taxes and benefits at smoothing shocks varies by the type of shock. As mentioned above, mortality and migration shocks are most effectively smoothed in the U.S. system when only tax adjustments are used, productivity shocks are most effectively smoothed when only benefit adjustments are used, and the variance for fertility shocks is lowest when adjustment relies equally on tax and benefit adjustments. Third, the ability of social security to smooth shocks of a given magnitude varies considerably across the types of shocks,

with a substantial proportional reduction possible for the effects of the mortality shock but little reduction for productivity, consistent with the intuition discussed above. Fourth, even where a reduction in variance is possible, some combinations of tax and benefit adjustments fail to smooth shocks. For example, while reliance primarily on tax adjustments lessens variance under the migration shock, variance increases when the primary adjustment is through benefits. Finally, considering the performance of the German and Swedish systems, we observe that the German system performs quite well relative to the U.S. system, with actually a lower variance than the lowest U.S. variant for both productivity and fertility shocks, even though the best-performing U.S. variant differs for these two shocks. The Swedish system, on the other hand, performs rather poorly, except in the case of the fertility shock.

To get a sense of what may underlie these results, we consider the effects of these shocks on different cohorts. Figures 6 and 7, respectively, display wealth equivalents and payroll tax rates for the different systems for each of the four shocks. (Note that the vertical scales in each figure differ across shocks.) For the U.S. system, we evaluate three variants, with only taxes adjusted, only benefits adjusted, and equal adjustment of the two.

a. Productivity Shock

As one would expect, a positive productivity shock increases well-being (Figure 6, top left panel), with the largest effects being experienced by the generations reaching adulthood around the time that the productivity shock hits. While the effects differ slightly across the different social security systems, these differences are very minor compared to the common impact of the shock itself.

All five systems appear to concentrate gains slightly more among cohorts reaching adulthood around the time the productivity shock hits. This makes sense given that unfunded

social security systems provide a rate of return based on the economic growth rate. The productivity shock makes social security a “less bad” deal, a source of gain that is absent when there is no social security. The U.S. tax adjust system appears to shift the gains slightly toward younger generations. Insight into the reason for this relationship comes from the top left panel of Figure 7, which shows the time path of the social security tax rate under the different systems. The tax rate is constant, by assumption, under the U.S. benefit adjust and Swedish systems. Under the U.S. tax adjust system, the tax rate drops initially, since it is easier to finance social security benefits with a more productive work force, and then gradually rises back to its initial level. Thus, younger workers get an added benefit through lower payroll taxes than in the other systems. The German system, on the other hand, appears to help the elderly slightly more and young workers relatively less. Unlike in the U.S. tax adjust system (and the U.S. both adjust system), the German system’s payroll tax rate overshoots its long-run level, because the productivity shock has dissipated but left in its wake an impact on social security benefits, which are based on lifetime earnings. That is, workers whose productivity has reverted to its original level must pay elevated benefits to those who worked during a higher-productivity period.

As indicated, the various social security systems appear to provide little, if any, cushioning of the productivity shock, and actually increase the benefit among those already most helped by the shock, for systems that use tax adjustments. This is consistent with our finding in Figure 5 that the systems provide little, if any, smoothing of the productivity shock, and that systems relying on benefit adjustment tend to perform better.

b. Mortality Shock

To interpret the wealth effects for the mortality shock, it is important to remember that they assess the change in the bundle of consumption and leisure using pre-shock mortality

profiles. By this measure, we would expect a shortening of lifespan to increase well-being, *ceteris paribus*, because it would make more resources available during the period in which an individual is alive.²³ This outcome is quite evident in Figure 6 for the no-social security case, in which older generations – those who are primarily affected by the temporary increase in mortality, experience an increase in welfare.

Why do those reaching adulthood just into the transition experience a small decline in welfare, at least in the no-social-security case? Wages dip temporarily, because older generations have less reason to save for old age and therefore accumulate less capital. Hence, those who reach adulthood shortly after the transition begins, who themselves will be largely unaffected by the mortality shock, experience lower wages and hence lower welfare. The same general equilibrium effects help older generations further, through a temporary rise in interest rates.

This pattern of effects across generations is substantially modified under the different social security systems, being softened under those that involve at least some adjustment of social security taxes. Under the U.S. tax adjust system, younger adults actually now gain because of the shock, while older generations see their gains reduced. The reason for this shift is the reduced payroll tax. A mortality shock temporarily reduces the old-age dependency ratio and hence allows a reduction in payroll taxes. (Unlike the case of the productivity shock, there is no subsequent need for a payroll tax increase under the U.S. tax-adjust system.) Thus, a social security system that incorporates payroll tax adjustments spreads the gains from a mortality shock, offsetting the negative general equilibrium effects on young workers' wages and also distributing to them some of the surplus made possible by the lower consumption needs of the

²³ Put another way, the pre-shock mortality profile assigns more weight to future periods of life than is consistent with the resources individuals must put aside for those periods, given actual life expectancy.

elderly. This smoothing is present to a lesser extent under the German system and the mixed U.S. system as well.

With only benefits adjusted, there is little smoothing, or even a reinforcement of the underlying effects of the shocks. The Swedish system, which incorporates survival probabilities in its annuity calculation, rewards those retiring around the time of the shock with higher annuity payments, reinforcing the annual consumption increase permitted by a faster decumulation of private assets. But this increase in benefits must be paid for by subsequent generations in the form of lower benefits, causing young workers to be even worse off than without social security.

Unlike in the case of productivity shocks, some of the social security systems have a large smoothing effect, as measured by their impact on wealth equivalents, although adjusting only benefits is much less effective for smoothing mortality shocks.

c. Fertility Shock

Recall that we measure welfare effects for fixed fertility profiles. Ignoring general equilibrium effects, we would expect a reduction in well-being for young adults alive at the time of the shock, because these adults must commit more resources to children as a result of larger family sizes. (As noted in footnote 20, this ignores the well-being generated by the voluntary decision to have more children.) That is indeed what occurs for the simulation in which there is no social security system. Under that scenario, there are essentially no “winners” from the boom in fertility. The biggest losers are those hitting adulthood at the time of the shock, who experience the shock’s full magnitude and its effect on their available resources per capita. But older working generations also lose, as young adults’ increased labor supply reduces their wages, and younger generations lose, as those hit by the shock are less able to accumulate capital.

With social security systems in place, the incidence of the shocks changes, even though the largest impact is still felt by the initial adult cohort. The patterns of incidence differ across the systems according to cohort. For cohorts in adulthood as of the shock, there is a small loss in welfare in the absence of social security, which may be traced to the reduction in wages that follows, with a delay, after the shock, due to the increased labor force. These cohorts are helped by systems that adjust benefits, notably the U.S. benefit adjust scheme, presumably because the improved system health generated by the lower old-age dependency ratio leads to an increase in their retirement benefits.

On the other hand, those entering adulthood after the shock are helped more by the U.S. tax adjust scheme than the US benefit adjust scheme, which makes sense because these individuals hit retirement only after the beneficial effects of the baby boom on the old-age dependency ratio have been reversed, while the tax adjustments lower their payroll tax rates during working years. This balance between the effects of adjusting taxes and adjusting benefits helps explain why the lowest-variance U.S. system (as shown in Figure 5) is the one that includes both adjustments.

Perhaps more surprising is that the Swedish system, which is the most beneficial to even younger cohorts, with a very long-lasting effect, and has the lowest variance of any of the systems. But this outcome may be understood once one takes into account that the benefit adjustments in the Swedish system are not based on achieving cash-flow balance. Individuals working when the system is healthy (because of rapid work-force growth) accumulate Notional Pension Wealth (NPW) at a faster rate. Thus, their eventual benefits may be enhanced, even if by the time they retire the growth of aggregate wages has slowed. Also, because increases in NPW have effects on benefits far into the future, the welfare of many cohorts may be affected.

d. Migration Shock

The relationship between immigration and the viability of PAYG pension systems has received ample attention in the literature (e.g., Auerbach and Oreopoulos, 1999, and Storesletten, 2000). But recent events, especially in Europe, provide a reminder that immigration levels may also be subject to considerable volatility. How effective are the different social security systems at spreading the economic shocks associated with sudden changes in immigration levels?

For the migration shock, the patterns of wealth equivalents are in some ways similar to those under a fertility shock, which is not surprising given that both involve additions to the population. However, while in the case of a fertility shock most generations do worse, with the migration shock many generations do better, particularly those generations reaching adulthood 30 to 40 years before the shock, who receive greater benefits when they retire, under the benefit-adjust Social Security policy. Unlike in the case of the fertility shock, there is no reduction in consumption due to child dependency, and therefore no related reduction in capital accumulation to fund consumption in retirement.

Comparing the two lower panels of Figure 6, we note that the effects of the migration shock die out more quickly, which is consistent with the shock being treated as a one-time event rather than one subject to an AR(1) process. In terms of the smoothing of shocks, perhaps the most noticeable difference between the two types of shocks is in the poorer performance of systems that adjust only benefits (the US benefit adjust and Swedish systems) under the migration shock. This appears due to the very sharp increase in benefits that follows immediately from a large influx of foreign workers, which leads to large increases in the wealth equivalents for workers nearing retirement. At the same time, younger workers are now part of a larger age cohort, which will put more pressure on their benefits when they retire.

e. Adjustments in the Swedish System

As our discussion highlights, the Swedish system, as we have modeled it, is unique among the systems we consider of being unfunded on a long-term basis but having financial asset accumulation and decumulation available due to buffer shocks. Figure 8 shows how the balance ratio, b , and the fund balance, F , relative to aggregate wages, react to the shocks we have considered. The balance ratio, in the upper panel of the figure, moves quickly to very near 1.0 following each of the four kinds of shock, but least so in the case of the fertility shock, where echo effects in population structure (see Figure 2) continue to cause small fluctuations.

The effects on financial assets (shown in the lower panel of the figure) are quite different under the four shocks. For a productivity shock, assets rise sharply in the short run, as only some of the additional tax revenues generated go to increased benefits. Assets then fall as the productivity shock decays and the impact of increased NPW accumulation (due to faster growth) eats into accumulated assets. Indeed, fund assets overshoot, becoming negative for a time until the slowing of benefit growth catches up with current economic fundamentals. For a fertility shock, the major impact on F is delayed, because the system is only directly affected when the larger cohort enters the work force. When it does, the effect is similar to that of a productivity shock, raising asset accumulation right away, before increased accumulation of NPW works to offset the asset accumulation. Unlike in the case of the productivity shock, however, the impact is more long-lasting and there is no overshooting, because of the echo effect of future large cohorts that keep the initial population shock from dissipating. For a mortality shock, there is little impact at all on trust fund accumulation, because changes in mortality above age 60 have no direct impact on the growth of wages, either through work-force growth (fertility) or wage growth (productivity). The small impact observed is due to the indirect effects associated with

labor supply and saving responses to the mortality shock. For a migration shock, the fund balance increase is gradual but not as delayed as under the fertility shock, because some migrants are already in the labor force and increasing payroll tax contributions. However, the fund balance peaks more rapidly as well, because of the shock's lack of permanence.

f. Summary

We have found that the various systems vary in their effectiveness at smoothing shocks, and also that as a group they smooth some types of shocks more than others. Given the differences in performance for different types of shocks, it is hard to say which system is best at smoothing shocks. However, we can consider overall performance by taking into account the relative importance of different shocks. For example, even though the smoothing of productivity shocks is minor, these shocks are by far the most important, in terms of the variance of wealth equivalents in the absence of social security; even minor smoothing of productivity shocks may be more important than more significant smoothing of, say, mortality shocks, which generate the smallest variance of wealth equivalents. Summing the variances of wealth equivalents under the four shocks for each system (as would make sense if each of the shocks occurred independently and with equal probability) yields results very similar to those for the productivity shock in Figure 5, with the important finding that (1) both the German and Swedish systems reduce the variance of wealth equivalents, but also that (2) the variants of the U.S. system that rely heavily on tax adjustment actually increase the overall variance of wealth equivalents.

4. Sensitivity Analysis

Our simulation model incorporates several simplifying assumptions and parameter choices. To illustrate the impact of such decisions, we consider the effects of two alternative

assumptions: varying one important parameter, the intertemporal elasticity of substitution, γ , and assuming that the economy is small and open, facing fixed world factor prices.

a. A Higher Intertemporal Elasticity of Substitution

We first consider a doubling of γ (from 0.35 to 0.7). For these simulations, we also adjust the pure rate of time preference (from -.015 to .0139) to preserve the same interest rate in the initial steady state without social security. The patterns are generally quite similar. However, there is a noticeable improvement in the performance of the Swedish system, in terms of the variance of wealth equivalents. This outcome appears to relate to the fact that the Swedish system adjusts faster when the intertemporal elasticity of substitution is higher. For example, the financial assets converge more quickly to their steady state value of zero.

With this improvement in the performance of the Swedish system, it now has an overall impact on the variance of wealth equivalents slightly better than the German system.

b. An Open Economy

It has long been understood that general equilibrium effects may moderate the effects of demographic change. For example, an increase in the old-age dependency ratio might increase the fiscal burden imposed on working generations, but these same generations could enjoy increased wage rates as a result of an increase in the capital-labor ratio. Here, we consider a different though related question, how the spreading of shocks depends on such general equilibrium effects. Of course, since all countries are open to a greater or lesser degree, it is also substantively important to understand how openness would alter the intergenerational effects of shocks under different pension systems. To address this question, we consider the effects of keeping returns to labor (in efficiency units) and capital fixed at their initial steady state values,

as might occur in a small open economy. Figure 9 shows the impact of this assumption on wealth equivalents for each type of shock, comparable to the base case results in Figure 6.

Comparing the results yields a number of interesting observations. For the productivity shock, there is a noticeable difference in the effects on generations reaching adulthood around 25 years into the transition, who do worse in the case of the open economy. In a closed economy, such cohorts enjoy higher wage rates, even though the productivity shock has largely eroded, because of a higher capital-labor ratio (generated by savings from higher past income). Thus, the gains from the productivity shock are generally more concentrated in the open economy. Also, the Swedish system now looks slightly better than the others, at least in terms of the spreading of the shock near its peak impact.

For a mortality shock, perhaps the most evident difference associated with the open economy is on cohorts hitting adulthood around the time of the shock. In a closed economy, these cohorts suffer a decline in wages, which can reduce their own lifetime welfare, at least for some of the social security systems. In an open economy, these cohorts do better, even in the absence of social security where they benefit from lower taxes needed to support other government expenditures. Here, the relative impacts of the different social security systems are similar to those in a closed economy.

The fertility shock has rather different effects on cohorts reaching adulthood between roughly 25 and 60 years after the shock, again doing better in the case of the open economy. Here, the mechanism is quite straightforward – a large population of workers no longer depresses the wage rate and therefore may actually do better, again because of lower taxes to pay for other programs (here because of a higher working population, rather than a lower dependent population). For this scenario, social security systems that rely on taxes do relatively worse in

the open economy than in the closed economy. In the case of the closed economy, the fertility shock that depresses wages also lessens the fiscal pressure on PAYG social security systems; adjusting payroll taxes reduces the burden on those who suffer wage declines, thereby cushioning the effects of the shocks. In the open economy, no such wage decline occurs, so reducing payroll taxes does not serve the same cushioning function.

Finally, for the migration shock, the patterns of wealth equivalents are quite different for cohorts around 0, i.e., those reaching adulthood at the time of the shock. Unlike in the case of a fertility shock, where the impact hits later cohorts, the effect here is immediate because the population shock involves current workers, not just future ones. As with the fertility shock, the affected working population does better in an open economy, benefiting from a broader population of taxpayers to cover general government spending but not suffering a decline in wages. As to social security, the relative performance of the different systems again is somewhat worse for systems that adjust taxes, but perhaps the more salient difference from the open economy is that all systems do worse in cushioning shocks. Put another way, without social security, the effects of a migration shock on the welfare of different cohorts are much smaller in an open economy, once the factor-price channel has been neutralized. All of the social security systems react to the migration shock either by increasing benefits – making older generations better off – or by reducing taxes – making younger generations better off – and these transfers exaggerate the underlying fluctuations in welfare associated with the shock itself, namely the effect of the shock via the general taxes and spending of government. As a consequence, the variance of wealth equivalents is lower in the absence of social security than under any of the alternative social security schemes.

In summary, excluding factor-price responses to the shocks considered here may reduce the desirability of using taxes rather than benefits to adjust PAYG social security systems, and may also reduce the benefits of social security as a mechanism for smoothing shocks.

5. Conclusions

We have simulated the ways in which different public pension structures spread the effects of isolated deterministic macro shocks across the leading and trailing generations in a general equilibrium setting. This enables us to make comparisons of outcomes across systems, across kinds of shocks, and across generations. The impulse-response approach pinpoints the particular consequences of different kinds of shocks, making it possible to observe and interpret these outcomes in a way that is not possible with stochastic simulations which show us the results of a mixture of kinds of shocks initiated at many different times.

The general equilibrium setting provides new insights about the effects of shocks filtered through different pension structures, which affect the way that shocks alter the saving and labor supply behavior of generations. For example, following a mortality shock, older working age generations have less need to save for retirement (life expectancy is shorter and annuity rates of return are higher, reflecting higher old age mortality), so capital per worker is reduced and wages fall, while the rate of return earned by the elderly on their assets rises. But these general equilibrium effects are modified in different ways by the pension structures, as discussed earlier.

Based on our measure of volatility, determined by variations in the wealth equivalents of shocks across generations, the pension systems we consider are effective to some degree in spreading shocks across generations, at least in some cases. However, the effectiveness depends on the type of shock as well as the choice of parameters. For some systems, parameterizations and shocks, social security systems actually concentrate the effects of shocks.

If we form an overall measure by summing the wealth equivalent variances arising from the four kinds of shocks, the most important of which is to productivity, we find that both the German and the Swedish systems reduce the variance of this overall measure, while some versions of the U.S. system that rely on tax adjustments can actually raise that variance. In an open economy the desirability of using taxes rather than benefits to adjust PAYG social security systems is reduced relative to the closed economy, and the open economy setting also reduces the benefits of social security as a mechanism for smoothing shocks – indeed a social security system can actually amplify the effect of some shocks in an open economy.

All of the systems considered here are unfunded systems, and yet they differ in their relative reliance on tax adjustments and benefit adjustments and on the extent to which they require annual pay-as-you-go balance rather than long-run balance (as under the Swedish system). Allowing a fluctuating buffer stock adds an additional tool for spreading risks across cohorts but, as our discussion of system stability highlights, also complicates the task of maintaining fiscal balance in an unfunded system. These findings suggest the need for further research regarding the design of unfunded systems that allow short-run fluctuations in annual pay-as-you-go balance.

References

- Ambler, Steven, and Alain Paquet (1994) "Stochastic Depreciation and the Business Cycle," *International Economic Review* 35(1), February, 101-116.
- Attanasio, Orazio, Andrea Bonfatti, Sagiri Kitao, and Guillermo Weber (2016) "Global Demographic Trends: Consumption, Saving, and International Capital Flows," *Handbook of the Economics of Population Aging*, vol. 1, 179-235.
- Auerbach, Alan J. (2002) "The Bush Tax Cut and National Saving," *National Tax Journal*, September, 387-407.
- Auerbach, Alan J. and Laurence J. Kotlikoff (1987) *Dynamic Fiscal Policy*. Cambridge University Press: Cambridge.
- Auerbach, Alan J., Laurence J. Kotlikoff, Robert Hagemann and Giuseppe Nicoletti (1989) "The Dynamics of an Aging Population: The Case of Four OECD Countries," *OECD Economic Studies*, Spring, 97-130.
- Auerbach, Alan J., Laurence J. Kotlikoff and Darryl R. Koehler (2016) "U.S. Inequality, Fiscal Progressivity, and Work Disincentives: An Intragenerational Accounting," NBER Working Paper 22032, February.
- Auerbach, Alan J. and Ronald D. Lee (2009) "Notional Defined Contribution Pension Systems in a Stochastic Context: Design and Stability," in J. Brown, J. Liebman, and D. Wise, eds., *Social Security Policy in a Changing Environment*. University of Chicago Press: Chicago, 43-68.
- Auerbach, Alan J. and Ronald D. Lee (2011) "Welfare and Generational Equity in Sustainable Unfunded Pension Systems," *Journal of Public Economics* 95(1-2), February, 16-27.
- Auerbach, Alan J. and Philip Oreopoulos (1999) "Analyzing the Fiscal Impact of U.S. Immigration," *American Economic Review* 89(2), 176-180.
- Bohn, Henning (2009) "Intergenerational Risk Sharing and Fiscal Policy," *Journal of Monetary Economics*, 56(6), 805-816.
- Diamond, Peter A. (1977) "A Framework for Social Security Analysis," *Journal of Public Economics* 8(3), 275-298.
- Feldstein, Martin S. and Jeffrey B. Liebman (2002) "Social Security," *Handbook of Public Economics*, vol. 4, 2245-324.
- Hasanhodzic, Jasmina, and Laurence J. Kotlikoff (2013) "Generational Risk-Is it a Big Deal?: Simulating an 80-period OLG Model with Aggregate Shocks." NBER Working Paper No.19179, June.
- Hong, Jay H., and José-Víctor Ríos-Rull (2007) "Social Security, Life Insurance and Annuities for Families," *Journal of Monetary Economics* 54(1), February, 118-140.

- Krueger, Dirk and Felix Kubler (2006) "Pareto-Improving Social Security Reform When Financial Markets are Incomplete!?", *American Economic Review* 96(3), 737-755.
- Ludwig, Alexander, and Michael Reiter (2010) "Sharing Demographic Risk - Who Is Afraid of the Baby Bust?," *American Economic Journal: Economic Policy* 2(4), 83-118.
- Matsen, Egil, and Øystein Thøgersen (2004) "Designing Social security – A Portfolio Choice Approach," *European Economic Review* 48(4), 883-904.
- Ríos-Rull, José-Víctor (2001) "Population Changes and Capital Accumulation: The Aging of the Baby Boom," *Advances in Macroeconomics* 1(1), 1-48.
- Nishiyama, Shinichi, and Kent Smetters (2007) "Does Social Security Privatization Produce Efficiency Gains?" *Quarterly Journal of Economics* 122(4), November, 1677-1719.
- Shiller, Robert J. (1999) "Social Security and Institutions for Intergenerational, Intragenerational, and International Risk-Sharing," *Carnegie-Rochester Conference Series on Public Policy*, vol. 50, 165-204.
- Storesletten, Kjetil (2000) "Sustaining Fiscal Policy through Immigration," *Journal of Political Economy*, 108(2), April, 300-323.

Figure 1. Baseline Fertility and Mortality Profiles
(probability of dying or giving birth in a year)

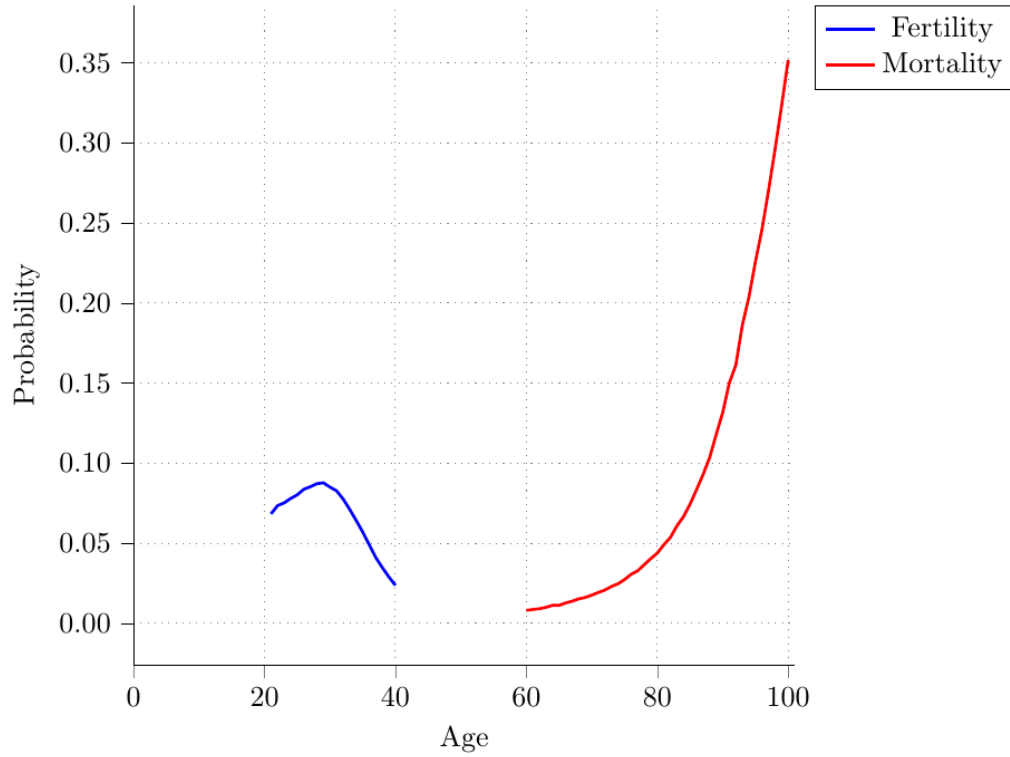


Figure 2. Population by Age, Fertility Shock
(adjusted for trend population growth)

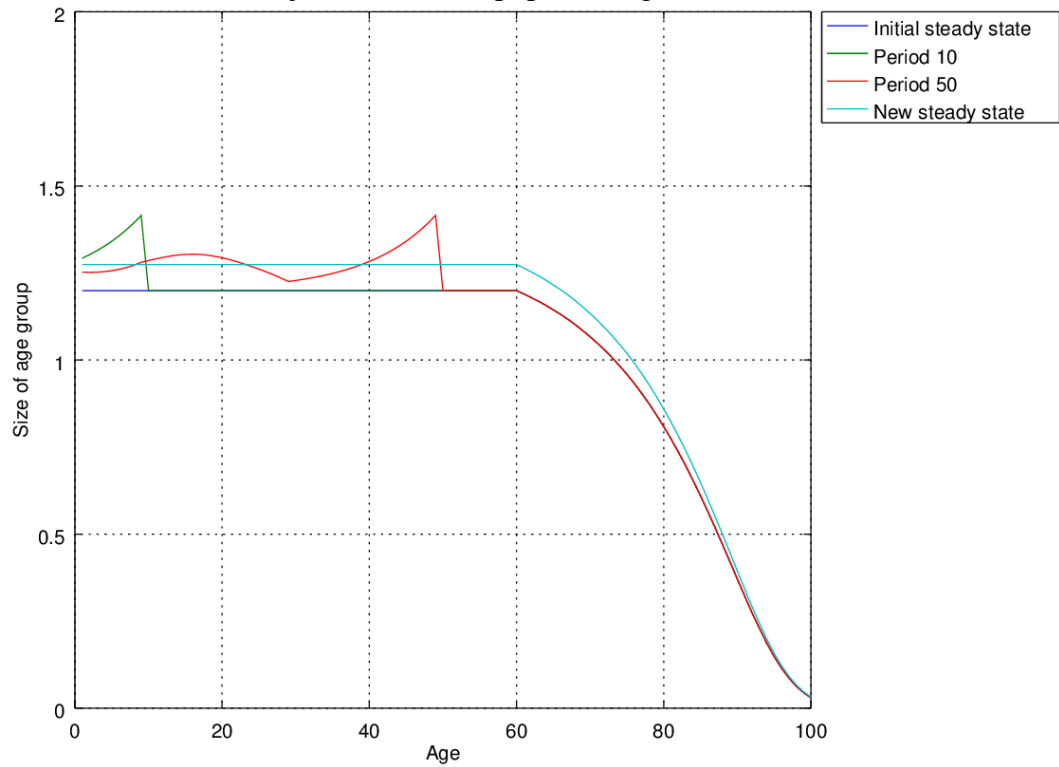


Figure 3. Aggregate Population (Normalized by Initial Population), Various Shocks

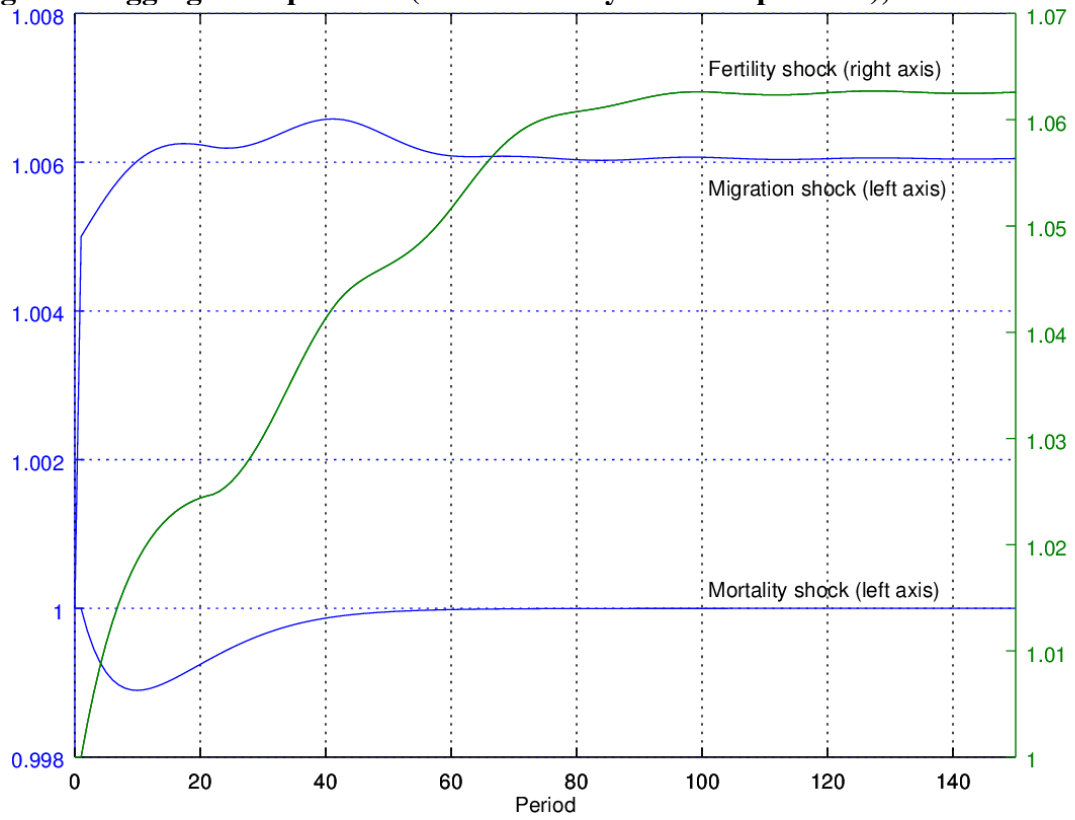


Figure 4. Migration: Age Profile of Immigrant Flow (to Germany in 2014)

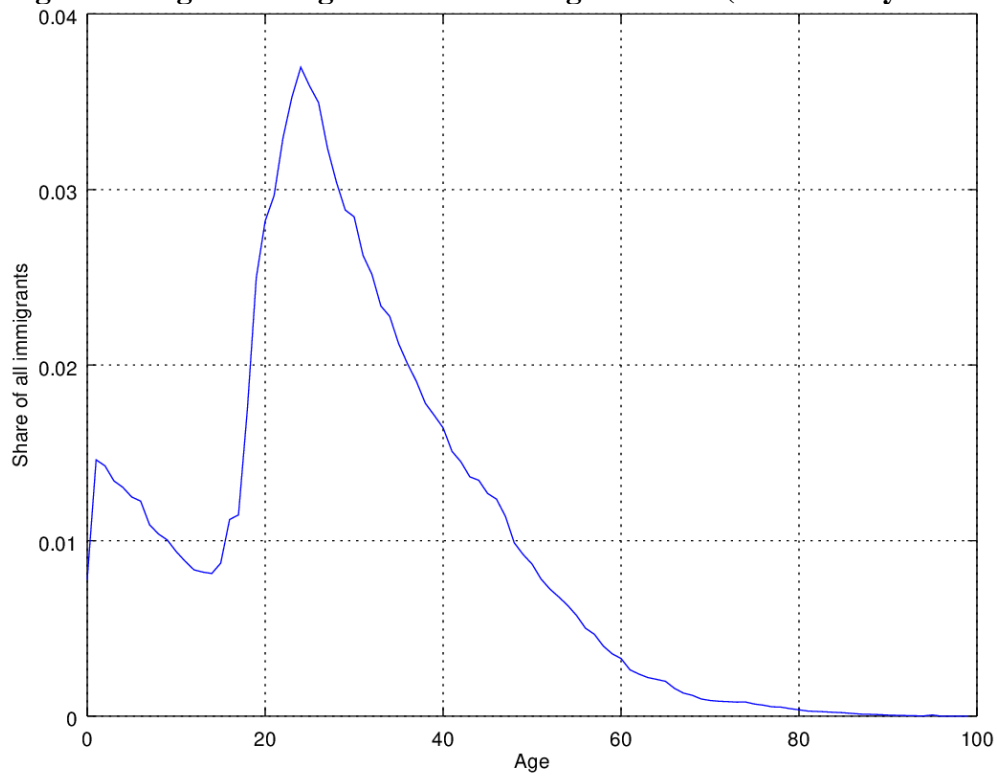


Figure 5. Wealth Equivalent Variances under Different Social Security Systems

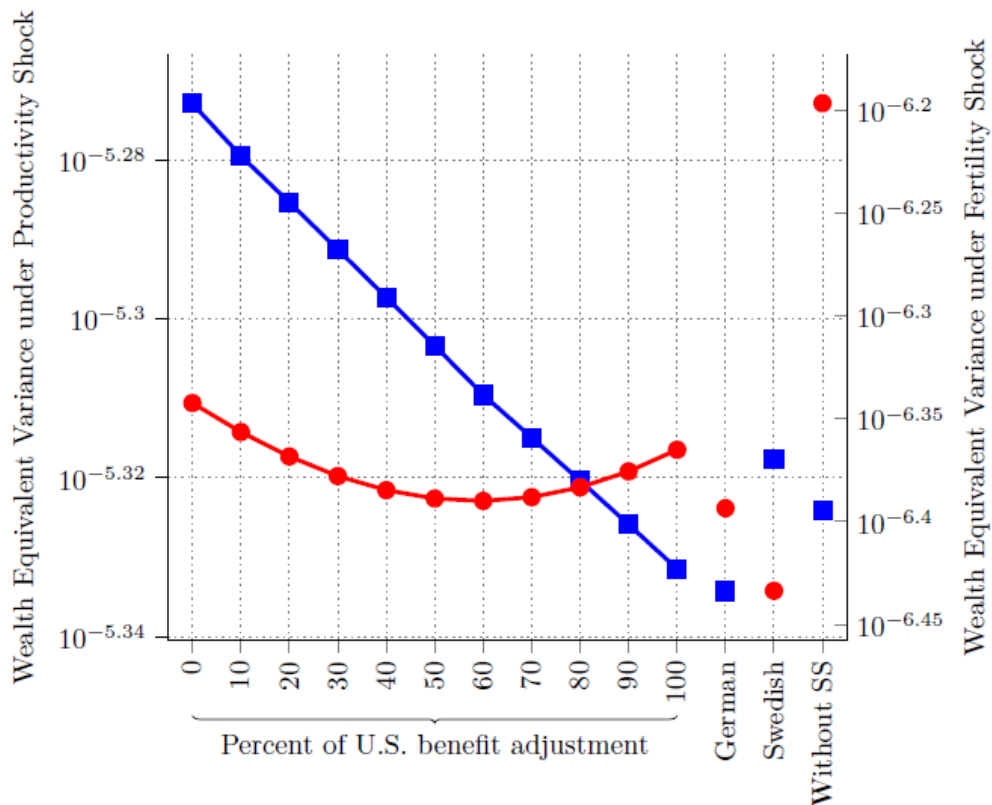
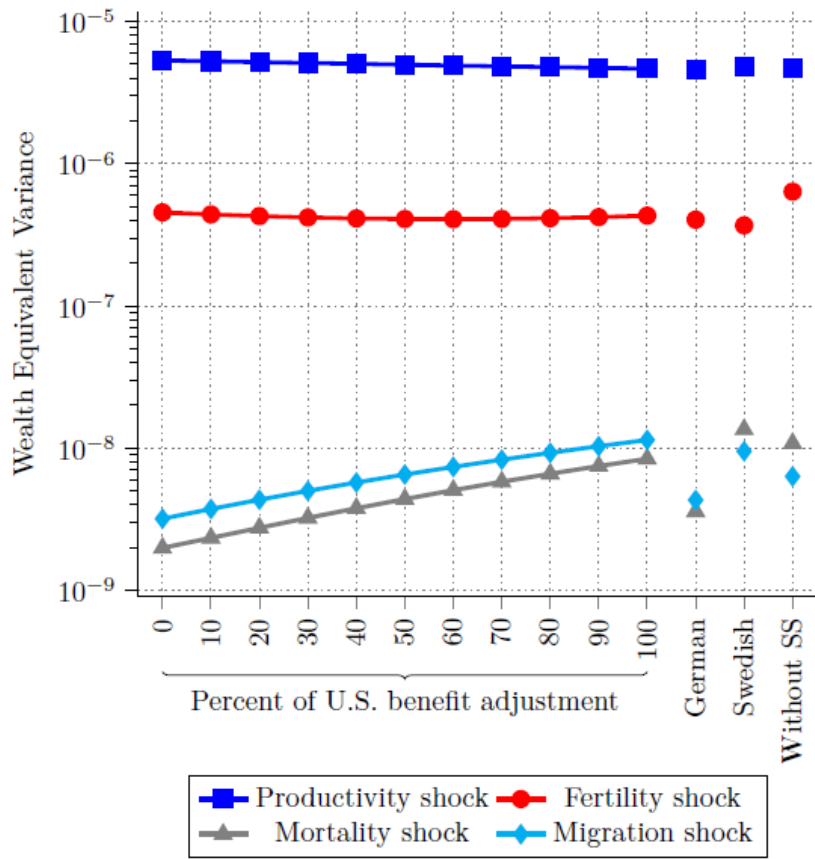


Figure 6: Wealth Equivalents by Year of Adulthood

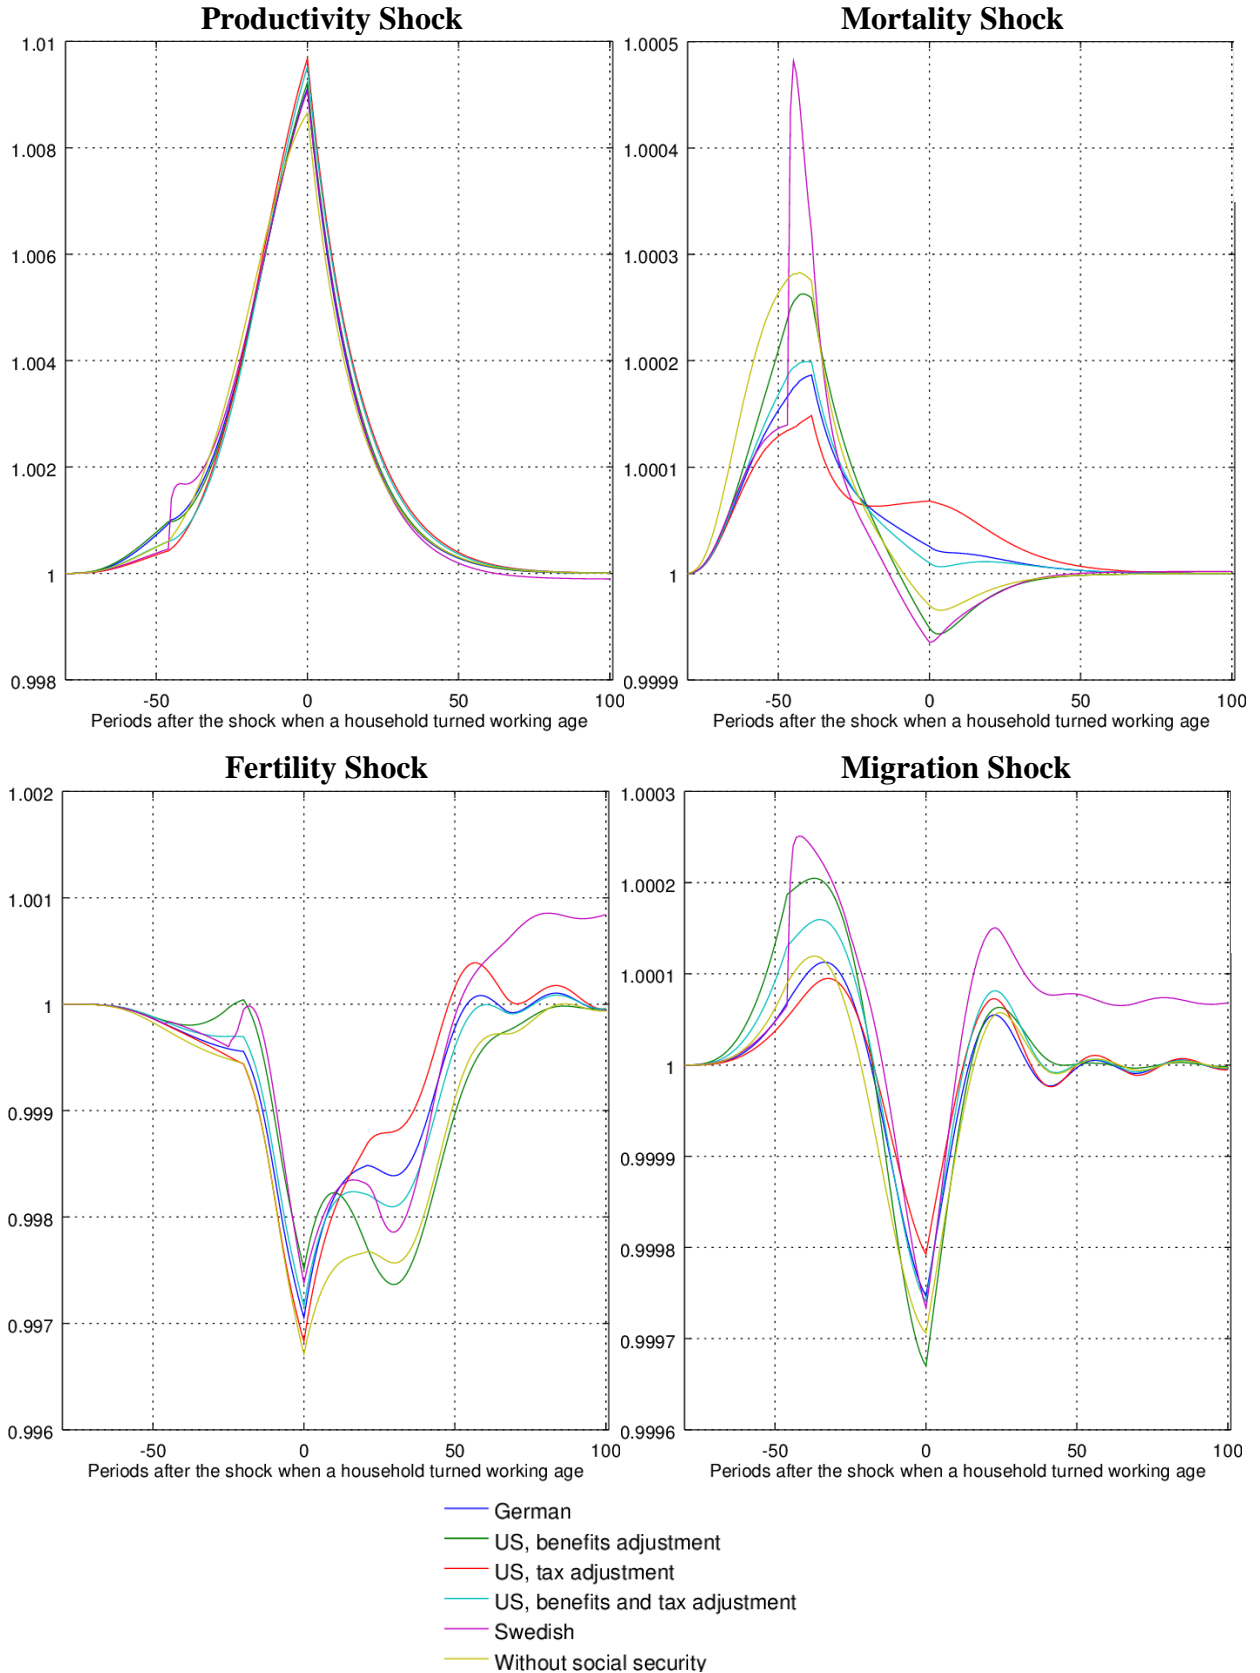


Figure 7: Social Security Tax Rates by Year

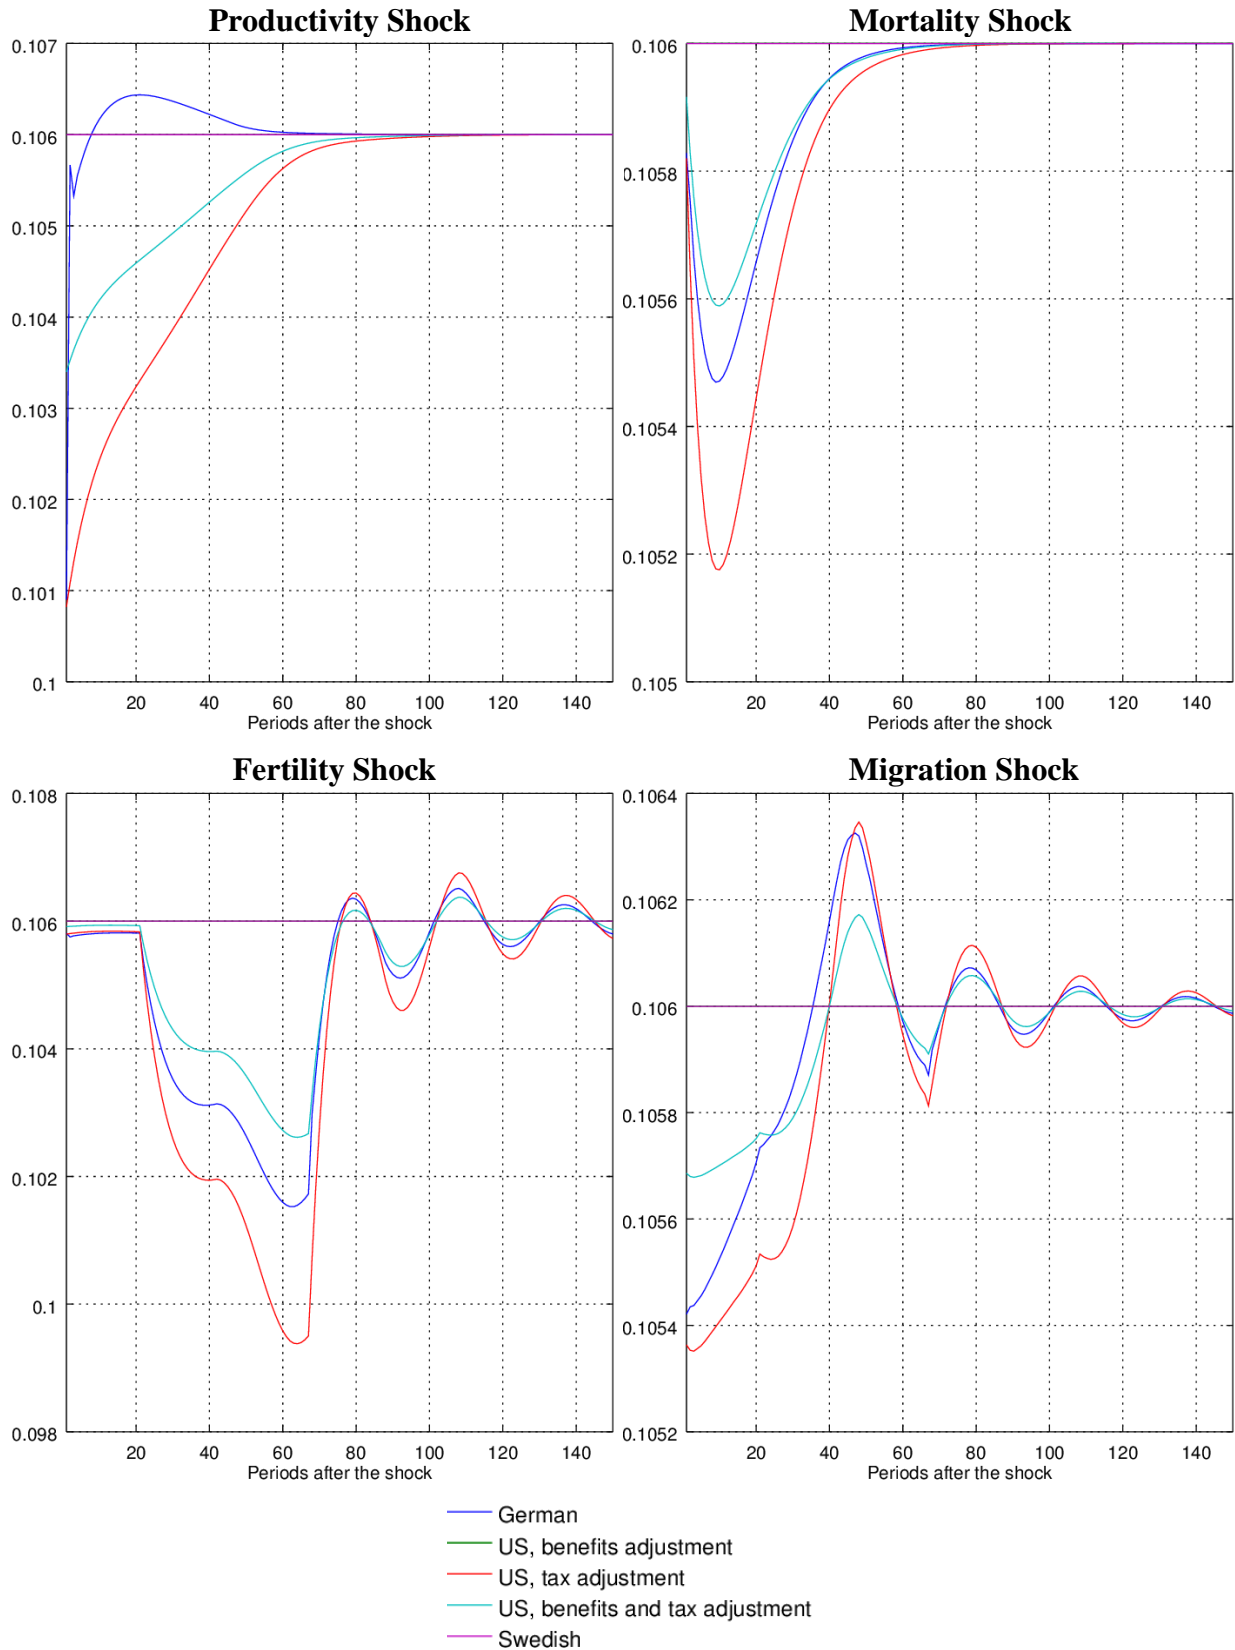


Figure 8a. Balance Ratio, Swedish System

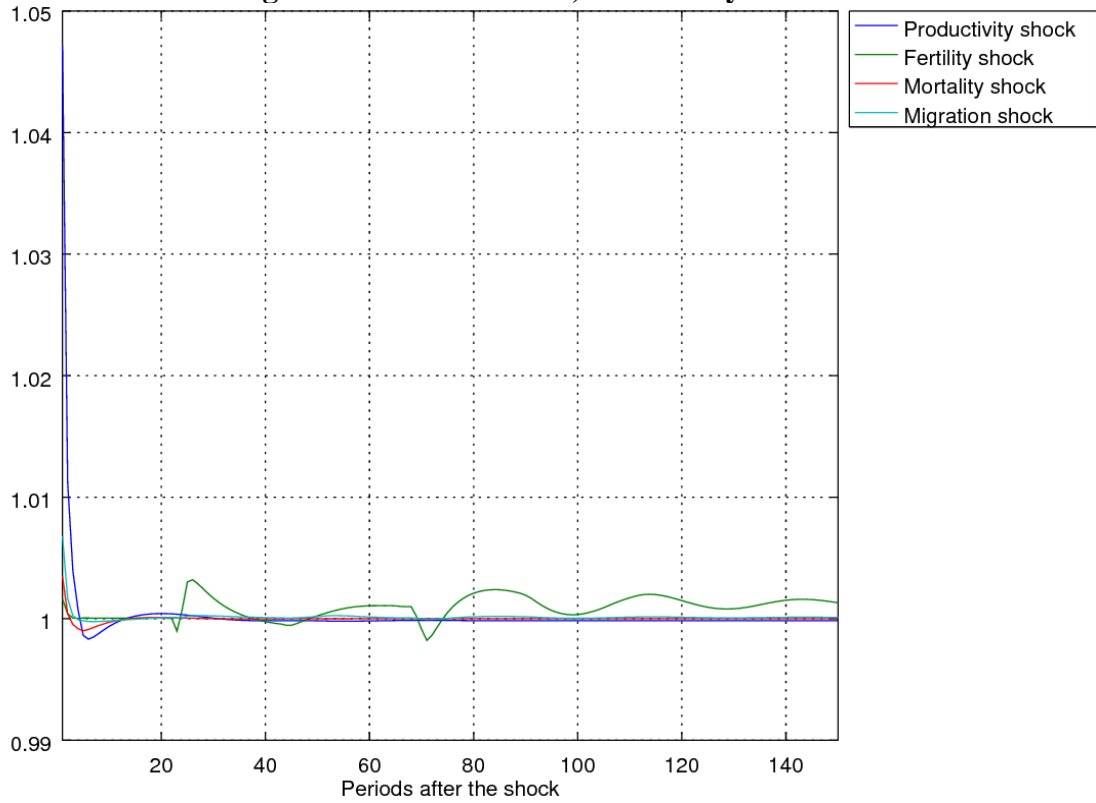


Figure 8b. Fund-Wages Ratio, Swedish System

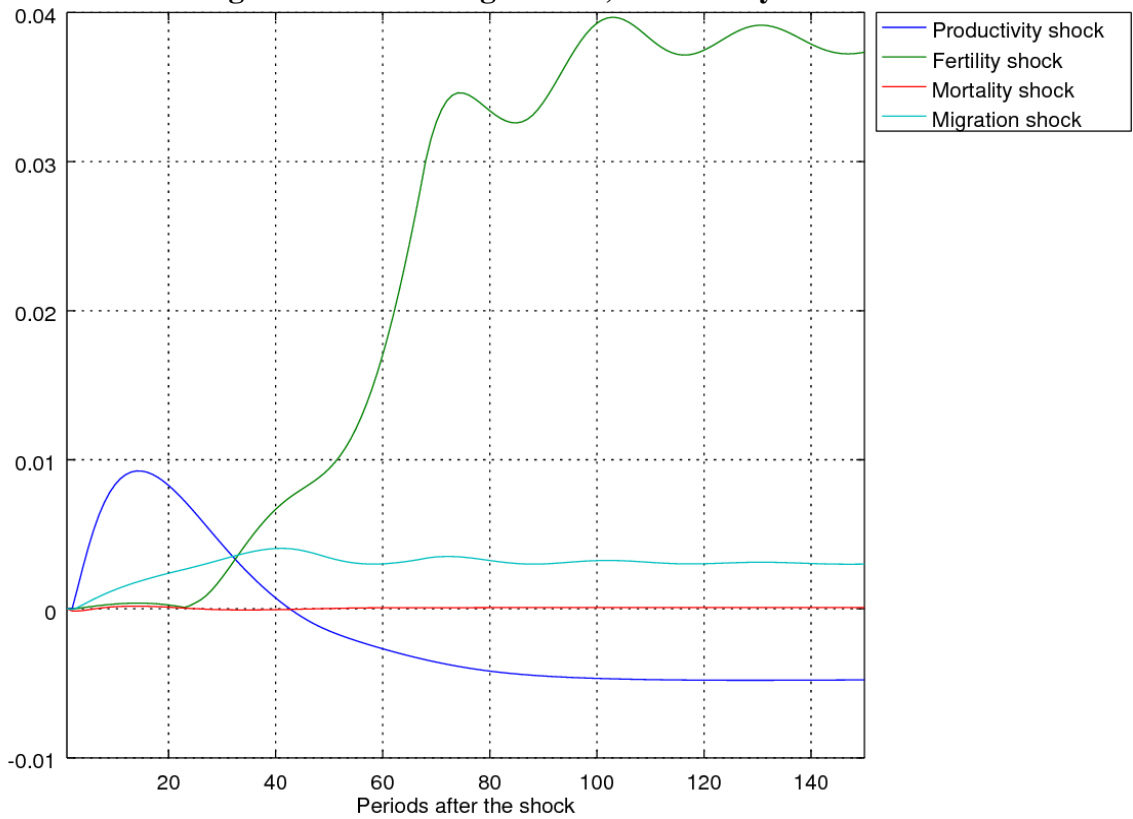


Figure 9: Wealth Equivalents by Year of Adulthood, Open Economy

