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, October 13, 2016

Abstract

Even with well-developed capital markets, there is no private market mechanism for trading between current and future generations, so a potential role for public old-age pension systems is to spread economic and demographic shocks among different generations. This paper evaluates the smoothing and propagation of shocks of five public pension schemes, based on the actual U.S., German and Swedish systems, which vary in the extent to which they rely on tax or benefit adjustments, and a trust fund buffer. Modifying the Auerbach-Kotlikoff (1987) dynamic general-equilibrium overlapping generations model to incorporate realistic patterns of fertility and mortality and shocks to productivity, fertility and mortality, we evaluate the effectiveness of these public pension systems at spreading the effects of such shocks, as well as shocks to migration. Our results suggest that system design, the source of shocks, and characteristics of preferences and technology are important factors in determining the potential of public pension arrangements to spread the burden of 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. 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.

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.

Two earlier papers (Auerbach and Lee 2009, 2011) studied a variety of existing and hypothetical unfunded arrangements, some adhering strictly to a PAYG approach and others allowing small fluctuations in trust fund balances. The first of these papers evaluated the stability of the existing Swedish system and several variants, while the second considered the performance of several stable unfunded systems, including the actual and hypothetical Swedish systems, the actual German system, and three stable variants of the existing U.S. social security system, 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.

Understanding the effects of an existing or proposed system on welfare is, ultimately, our objective in studying how different shocks are spread among generations, but looking more closely at the pattern can help us understand why certain systems seem to perform better in the welfare dimension and also how systems different than those we have considered would perform in response to different patterns of shocks. That is, our past welfare analysis was based on empirically estimated demographic and economic stochastic processes for the United States, but patterns in other countries, or in the United States in the future, may differ, and it would be useful to have a more general picture of how different systems perform in response to different types and patterns of shocks. Therefore, in this paper, we develop a methodology for isolating the effects of different types of shocks on the welfare of different generations, looking in particular at the extent to which such shocks are spread across cohorts.

Previous Research

As our earlier work provides a foundation for the current project, we give a brief summary of its major features and findings here.

Following Lee and Tuljapurkar (1994), we incorporated estimated stochastic processes for fertility and mortality in the United States along with an assumed deterministic immigration level to generate a stochastic population process. The real rate of return and the rate of labor productivity growth were also modeled as stochastic time series, with the long run mean values of these stochastic processes constrained to equal the central assumptions of recent Trustees' projections (from 2004). We assumed that wages tracked labor productivity. In order to make the study of long-term patterns easier, we adopted two strategies. First, we started our stochastic simulation from initial population and economic conditions but threw out the first hundred years to eliminate the significance of initial conditions. Second, we modified the stochastic processes to remove drift terms. That is, we wished to study a stochastic equilibrium in which the means or expected values of fertility, mortality, immigration, productivity growth, and interest rates had no trend, and the population age distribution was stochastically stable rather than reflecting peculiarities of the initial conditions.[1] To determine the distributions of outcomes, we drew 1000 random trajectories, each for a period of 500 years after the initial 100-year period. We used these trajectories as a platform for studying different social security systems.

We limited our consideration to social security systems that were stable, in the sense of not being subject to excessive debt-payroll ratios on some trajectories. Systems that failed this stability test would require future policy changes, and so it would not have been meaningful to consider how they spread shocks without incorporating these policy feedbacks. The systems

[1] The resulting population processes had some trend in variances if not mean, but our simulation experiments showed these trends to be small within the horizons we used, over the range one hundred to six hundred years.

analyzed in Auerbach and Lee (2011), who considered the welfare properties of each, calibrated to US economic and demographic data, included the following:

1. The existing Swedish system, which uses Notional Defined Contribution (NDC) accounts, a system that bases “notional pension wealth” accumulation and ultimate annuitization at retirement on the rate of productivity growth and includes an automatic balancing mechanism that can come into effect to ensure stability.
2. Three variants of the Swedish system developed in Auerbach and Lee (2009) that use different versions of the balancing mechanism (the “brake”) and/or incorporate labor force growth as well as productivity growth in the calculation of pensions.
3. The German system, which bases pension growth on productivity growth and the growth rate of the old-age dependency ratio and uses payroll taxes as a residual to accomplish annual PAYG balance.
4. Three variants of the US system that achieve fiscal stability through the use of uniform tax and benefit adjustments that accomplish PAYG balance, as illustrations of what the US system might look like if it were financially stable.

New Analysis

Under five of the public pension systems just discussed (one variant of the Swedish system, the German system, and the three variants of the US system), we wish to consider how shocks of particular types play out over time and generations, using estimated impulse responses to shocks. One initial thought might be to carry out such analysis using the stochastic modeling approach of our previous work by looking only at the particular shocks of interest, one at a time. However, such an approach is difficult, because each type of shock has complex economic effects and channels that cannot be determined without an explicit general equilibrium model.

For example, a fertility shock would affect not only the relative sizes of different cohorts, and hence the finances of a public pension system, but also the returns to labor and capital over time. Thus, to determine how a particular pension system spreads the shocks arising from a fertility shock, we need a full general equilibrium model to trace through all of these effects.

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, and assumed a smooth rate of productivity growth. We will go through the model, indicating the adaptations developed for use in this paper and how we use the newly modified model to measure the effects of shocks to productivity, fertility and mortality.

The Model

The model we use is one in which individuals live for up to *T* years, the first 20 of which are spent as minor dependents of parents who make consumption decisions on their behalf. At any given time, the household consists of one parent and minor children. Household utility in each year is based on the parent's consumption and leisure, following a CES function, and each child's consumption. The household maximizes 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 fertility and mortality experience, with deterministic overall cohort fertility and mortality profiles of probabilities of experiencing mortality or fertility.

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.[2] We assume that births 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 $s_{at} = (1\text{-}\ m_a\mu_t)$.[3] 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

[2] 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% of female births die before age 60, and of those who survive to begin childbearing at age 20, only 8% 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%.

[3] A more general specification for the survival probability would be easy to include in the model. The same is true for the fertility specification.

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 the baseline fertility and mortality hazard rates we use. These are for the United States, taken from the Human Mortality Database and the Human Fertility Database for 2010,[4] 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.[5]

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:

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

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_j 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., $\omega_j = .25 + .25*(j-1)/19$. In expression (1), there are also four household preference parameters: δ is the pure

[4] The data are available at http://www.mortality.org/ and http://www.humanfertility.org/cgi-bin/main.php.

[5] 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 approximately the correct measure of life expectancy at age 60, although overstating life expectancy at earlier ages.

rate of time preference, α is the leisure intensity parameter, ρ is the intratemporal elasticity of substitution between consumption and leisure, and γ is the intertemporal 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 US before-tax rate of return to capital measured in Auerbach, Kotlikoff and Koelher (2016), just under 6.5 percent. The resulting value for δ is -.015.[6]

The economy has one production sector, in which the representative firm is assumed to behave competitively in factor and output markets and produce 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 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 date of retirement being endogenously determined by preferences and factor prices. This date of retirement may vary over time and is distinct from the initial age at which benefits are received, which we set at 67 for all social security systems considered, consistent with the normal retirement age under the current US 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:

[6] This low time preference rate may be rationalized as offsetting the model's lack of two motives that would increase private saving: precautionary saving (which would be present without a complete annuities market) and bequests.

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 the consumption of 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 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.[7]

Finally, we include the government sector in the model. The government sector consists 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 productivity growth) relative to their respective population groups and non-age-based components (e.g., defense spending) held constant (except for productivity 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, and o (young, middle-aged, and old) or for the total population. That is, overall general government spending at any date t equals:

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

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

[7] While productivity shocks hit both interest rates 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.

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 over time except as a result of changes in the age structure of the population.

General government is funded with a proportional income tax and a consumption tax, and we ignore the use of government debt for the general government budget. In the initial steady state, we set the proportional income tax to 20 percent and the consumption tax equal to 3 percent, which are similar to recent estimates for the US tax system from Auerbach (2002). Because revenue requirements to meet the required spending, as specified in expression (2), fluctuate during the transition after a shock, we allow the consumption tax rate to vary to ensure annual budget balance for the general (non-pension) government.

Each public pension system is modeled on one of the various systems described above. In the results that follow, four of these five systems by design have year-by-year budget balance – the US "benefit adjust" system, in which the payroll tax is fixed and all adjustments occur to benefits, the US "tax adjust" system, in which the replacement rate is fixed and all adjustments occur to payroll taxes, the US "tax and benefit adjust" system (also referred to as "US both"), in which 50 percent of the cash-flow adjustment occurs on the tax side and 50 percent on the benefit side, and the German system, which incorporates adjustments to both taxes and benefits in any given year. The remaining system – a variant of the Swedish system – has a trust fund that serves as a buffer stock; only benefits are adjusted, according to a balancing mechanism that increases or decreases the benefit growth rate in response to a funding ratio based on the trust fund balance and various measures of system implicit assets and liabilities.

For any version of the US 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.[8] We then solve for the payroll tax such that budget balance is achieved, according to the expression,

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

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, according to whether we are considering the benefit adjust scenario, the tax adjust scenario, or the scenario in which benefits and taxes adjust.

For the German system, 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

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

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

$$(5) \qquad B_t = B_{t-1} * \frac{w_{t-1}(1-\theta_{t-1})}{w_{t-2}(1-\theta_{t-2})} * \left[1 - .25 * \left(\frac{OA_{t-1} - OA_{t-2}}{OA_{t-2}}\right)\right]$$

[8] 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.

where w_t is the aggregate wage in year t, normalized for age-specific productivity differences, and OA_t is the old-age dependency ratio in year t, which we define to be the ratio of the adult population in retirement (age > 66) to those not in retirement (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 US 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 US 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.

The Swedish system is a notional defined contribution (NDC) system, in which each worker accumulates notional pension wealth in a virtual account and then receives an annuity based on *notional pension wealth* 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 NDC 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:

$$(6) \qquad NPW_{it} = NPW_{it-1}*(1+r_{it-1}) + T_{it-1}$$

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.[9]

Upon reaching the retirement age, R, which we will mark by the exogenous age of benefit receipt used in all of our systems[10], 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:

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

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.

[9] The actual Swedish system is based on wages per worker, but with a growing or declining labor force this is not consistent with a long-run balance between taxes and benefits.

[10] Note that the age of benefit receipt is independent of labor force participation, as is true in the United States for individuals who have reached the normal retirement age (although such individuals can choose to defer benefit receipt until age 70). We also assume that once individuals begin receiving benefits they no longer are subject to the payroll tax, even if they continue working.

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

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

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), 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 $(g_t - \bar{g})/(1+\bar{g})$.

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. For the simulations reported below, we set the tax-benefit linkage at 0.25, meaning that one fourth of payroll taxes are ignored when individuals make labor supply decisions.[11]

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 three types of shocks in year 2, with no further shocks thereafter (but the shock itself 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. Such changes do not occur in the U.S. tax adjust system, where the benefit rate is fixed at its initial level, the U.S. benefit adjust system, where the tax rate is fixed at its initial level, or the US both adjust system, where adjustments are tied to both initial levels. Our variant of the Swedish system also returns to its original state because of the balancing mechanism, which 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

[11] One might expect there to be differences in linkage among the systems. For example, one of the arguments in favor of notional defined contribution systems like Sweden's is greater transparency in the linkage between taxes and benefits. However, we lack evidence on the variation in linkage across the systems and do not anticipate that such differences would exert a major influence in the extent of smoothing of shocks.

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 discussed below, however, the post-transition tax rates in the German system are very close to the original one.

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.

Once the transition path is determined, we can then calculate how the welfare of each cohort is affected by each particular type of shock, where the household's welfare is based on its expected utility given in expression (1). From this, we calculate the *wealth equivalent* of the cohort's utility change, x, as the scalar that, when multiplied by the household's vector of consumption and leisure, equalizes steady state utility and actual utility along the economy's transition path in the presence of the shock. By analyzing how the effects of the different types of shocks are spread among different cohorts, we can gain insight into how and why the different public pension arrangements lead to different overall welfare when analyzed in the context of multiple shocks of all types, as in our previous analysis (Auerbach and Lee 2011).[12]

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

[12] Note that the levels of wealth equivalents are not directly comparable across different social security systems, because the initial equilibria in the absence of shocks are different for the different systems.

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.[13] 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. 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.[14] A complete smoothing of shocks would correspond to equal wealth equivalents for all individuals, spreading the costs or benefits of the shocks in proportion to initial resources.

Basic Results

For each of the shocks considered, to fertility, mortality, and productivity, we assume a one-time shock followed by decay based on an AR(1) process. Based on estimates using postwar US 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. Because the model is nonlinear, the

[13] 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.

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

results cannot simply be scaled down in proportion to determine the effects of smaller shocks, but we considered the effects of shocks one half as large and found little qualitative impact to the results discussed below.

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 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.

Figures 4 and 5 show 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 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 <u>and</u> level.

Stability of the Swedish System

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, one 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. But there is no assurance that this steady state is stable or unique. Indeed, for many reasonable parameter combinations, this equilibrium is unstable, and the only stable steady state is one for which $F > 0$ and $b > 1$.

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 (say, an increase in productivity) 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.

Note that this accumulation of trust fund assets is occurring in a variant of the Swedish system designed to work against positive asset accumulation, that is, with a symmetric balancing mechanism rather than the actual Swedish system's asymmetric mechanism. This suggests that there may be powerful forces in the direction of asset accumulation within the actual system, where the balancing mechanism is turned off 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.

The results presented below are for parameterizations where the "normal" steady state in the Swedish system is stable.[15] As a benchmark for each experiment, we consider the impact on the economy in the absence of any social security system.

[15] For our base case parameter assumptions, stability is just satisfied – a small increase in the pure rate of discount, δ, leads to instability.

Productivity Shock

Figures 6a-d show the effects of a productivity shock on the welfare of individual cohorts, the economy, and the social security system. Figure 6a shows the impact on well-being, based on the wealth-equivalent measure described above. As one would expect, a positive productivity shock increases well-being, 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. Figures 6b and 6c show the corresponding impacts over time on the aggregate wage rate (the wage rate for labor supply of unit efficiency, normalized to 1 in the initial steady state) and the interest rate. Not surprisingly, both jump up when the productivity shock hits, since both labor and capital become more productive. Again, differences across social security scenarios are minor, except for the interest rate being generally lower in the absence of social security (Figure 6c). This difference is due to the well-known negative effect of unfunded social security systems on capital accumulation; a lower capital stock leads to a lower capital-labor ratio and a higher interest rate.[16]

One interesting phenomenon in Figure 6c is that interest rates overshoot when adjusting back to their original level, falling below their initial value before recovering. The explanation for this is that, under perfect foresight, individuals who know that their wages are higher now than they will be in the future concentrate labor supply during the period of temporarily high productivity. In doing so, they accumulate a lot of capital to be used to finance higher levels of consumption later in life, which temporarily lifts the capital-labor ratio above its long run value after the productivity shock itself has dissipated.

[16] Note that there is no comparable impact on the wage rate in Figure 6b because the wage rate shown here is normalized to 1 in the initial steady state.

As to the welfare effects of the different social security systems (Figure 6a), if one compares the patterns of gains to those in the absence of social security, all four 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. Among the four social security systems, the US tax adjust system appears to shift the gains slightly toward younger generations. Some insight into the reason for this relationship comes from Figure 6d, which shows the time path of the social security tax rate under the different systems. The tax rate is constant, by assumption, under the US benefit adjust and Swedish systems. Under the US 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 US tax adjust system (and the US both adjust system), the payroll tax rate in the German system actually 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. One can get an overall sense of the extent of smoothing by considering each system's impact on the variance of the wealth equivalents. The lower panel of Figure 6a shows the ratio of the variance of wealth equivalents (over all generations born from 100 years before to 800 years after the shock) with social security to the variance without social security. Indeed, none of the systems have a large impact on the variance, although three of the systems actually increase the variance of wealth equivalents[17], with the systems that adjust both taxes and benefits – the US both adjust system and the German system – performing best by this measure.

Mortality Shock

Figures 7a-d show the effects of a mortality shock, beginning with the impact on individual welfare in Figure 7a. To interpret this figure, it is important to remember that it assesses 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.[18] This outcome is quite evident in the figure 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 shortly after the transition begins experience a small decline in welfare, at least in the no-social-security case? The explanation appears in Figure 7b, which shows the wage-rate trajectory over time. 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

[17] These variances are weighted by population. The relative results are similar when calculated without population weights.

[18] 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.

general equilibrium effects help older generations further, through a temporary rise in interest rates (Figure 7c).

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 US tax adjust system, younger adults actually now gain as a result of the shock, while older generations see their gains reduced. The reason for this shift is the reduced payroll tax, shown in Figure 7d. 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 US 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 elderly. This smoothing is present under the German system as well, which also is sufficient to reverse the welfare effect on young workers, and to a much lesser extent under the US both adjust system, with both effects being smaller than under the US tax adjust system because only part of the adjustment to the shock is through changes in the tax rate.

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. Now, the US tax adjust system, which performed worst for productivity shocks, performs best, although the two systems with tax and benefit adjustments (the US both adjust system and the German system), which showed the best performance under the productivity shocks, are nearly as effective in this case. Adjusting only benefits is much less effective for smoothing mortality shocks, and the Swedish system has virtually no effect on the variance.

Fertility Shock

Figures 8a-8d show the effects of a fertility shock. Again, recall that the welfare effects are measured 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. That is indeed what occurs in Figure 8a for the simulation in which there is no social security system. Under that scenario, there are essentially no "winners" from the boom in fertility, and 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.

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 sharp reduction in wages that follows, with a delay, after the shock (Figure 8b). These cohorts are helped by systems that adjust benefits – the US benefit adjust scheme and the Swedish system – 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 US 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 (Figure 8d). Perhaps more surprising is that the Swedish system, which also adjusts only benefits, is the most beneficial to even younger cohorts, and with a very long-lasting effect. 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.

As the lower panel of Figure 8a shows, all systems but the Swedish system successfully smooth fertility shocks, with quite similar overall effects. The Swedish system is by far the least effective, by this measure, substantially *increasing* the variance of outcomes.[19]

Adjustments in the Swedish System

As this last discussion highlights, the Swedish system, as we have modeled it, has the unique feature among the systems we consider of being unfunded on a long-term basis but having financial asset accumulation and decumulation available to buffer shocks. Figures 13 and

[19] The Swedish system appears to provide larger gains, or smaller losses, to most cohorts than some of the other systems, notably the US both adjust system. Recall, though, that, as discussed in footnote 12, values of wealth equivalents are not directly comparable across different systems, even for the same shock.

14 show how the balance ratio, *b*, and the fund balance, *F*, relative to aggregate wages, react to the different shocks we have considered. The effects on financial assets are quite different under the three shocks.

For a productivity shock, assets rise sharply in the short run, as the additional tax revenues generated are not completely used immediately to increase benefits or cut taxes. They 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, finally, 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.

An Extension: Migration Shocks

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?

While there is no immigration in our basic model – the population evolves solely through fertility and mortality of the domestic population, 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. Figure 9 shows the profile we use, based on immigrants to Germany in 2014.[20] 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.[21] Figure 10 shows the impact on population age structure and Figure 11 on total population, relative to the no-shock trajectory.[22] 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.[23]

Figures 12a-12d show the impact of the migration shock on wealth equivalents, wages, interest rates, and payroll taxes. The patterns of wealth equivalents have a similar shape to those under a fertility shock, which is not surprising given that both involve additions to the population. Comparing Figures 12a and 8a, 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

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

[21] https://www.destatis.de/EN/FactsFigures/SocietyState/Population/Migration/Tables/MigrationTotal.html.

[22] Note that, unlike in the case of a fertility shock, the population level overshoots in the case of the migration shock because some of the new population is above child-bearing age.

[23] Having different economic state variables than domestic residents of the same age, these immigrants choose different paths for consumption and labor supply, even though they are assumed to have the same labor productivity and preferences.

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.

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. For example, as just discussed, systems with benefit adjustments perform worse in smoothing migration shocks than for fertility shocks. And none of the systems are very effective in smoothing productivity shocks. Given the differences in performance for different types of shocks, it is hard to say which system is best at smoothing shocks, but some tentative conclusions are possible.

First, systems that include both tax and benefit adjustments – among those we have considered, the German and US both adjust systems – are more robust in dealing with different types of shocks. By comparison, the US tax adjust system performs very well in smoothing fertility and mortality shocks, but poorly in smoothing productivity shocks; and both the German and US both adjust systems reduce the variance of wealth equivalents more than the US benefit adjust and Swedish system for <u>every</u> shock considered.

Second, even where there is no strict ranking among systems, we can say something about relative 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 the most important, in terms of the variance of wealth equivalents in the absence of social security;

so 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 lowest sums for the German and US both systems, with the US tax adjust system and the Swedish system performing worse than no system at all, in the former case because of its exacerbation of the productivity shock.

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 Higher Intertemporal Elasticity of Substitution

Figures 15-18 show the impact of doubling γ (from 0.35 to 0.7) on wealth equivalents for each of the four types of shocks, repeating in the upper panel the results from Figures 6a, 7a, 8a, and 12a for $\gamma = 0.35$.[24] For these simulations, we also adjust the pure rate of time preference (from -.015 to .0139) to preserve the same initial interest rate in the initial steady state without social security.

Comparing the upper and lower panels of these figures, we can observe that the patterns are generally quite similar. However, there is a noticeable improvement in the performance of the Swedish system. This is perhaps most visible in Figure 17, for the fertility shock, but it is

[24] A value of $\gamma = 0.7$ is higher than has often been found in the literature, but there have been some estimates finding an even higher degree of intertemporal substitution. See, for example, Gourinchas and Parker (2002).

also the case for the productivity shock (Figure 15). 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 (as shown in Figure 14 for the base case) 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 similar to the German, US benefit, and US both systems. The US tax adjust system now shows the worst performance, although it also now produces a slightly lower variance than exists in the absence of social security.

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 increases in their 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. 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. Figures 19-22 show the impact of this assumption on wealth equivalents for each type of shock, again repeating in the upper panels the base case results from Figures 6a, 7a, 8a, and 12a.

Comparing the upper and lower panels of these four figures yields a number of interesting results. In Figure 19, 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). The effect of this higher capital-labor ratio can be seen in the overshooting of the interest rate in Figure 6c. 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, in Figure 20, 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 (Figure 7b), 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, as shown in Figure 21, 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 (Figure 8b) 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, shown in Figure 22, 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 shocks by either increasing benefits – making older generations better off – or 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.

In summary, excluding factor-price responses to the shocks considered here may affect the desirability of using taxes rather than benefits to adjust PAYG social security systems. The Swedish system's overall performance is now better than that of the US tax adjust system, as was the case for the higher value of γ, and the US benefit adjust system now performs slightly better than the German system, which relies more heavily on tax adjustments than the US both adjust system. However, the German system still performs better than the US tax adjust or Swedish systems, and the US both adjust system has the lowest overall variance of all, confirming the general attractiveness of a strategy that relies on both tax and benefit adjustments.

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. Under the tax adjust system, for example, the mortality shock reduces the number of retirees, leading to a lower tax rate for workers, offsetting the reduction in wages. Following a fertility shock, once the incremental births enter the labor force, wages fall and interest rates rise. The higher interest rates benefit the elderly retirees, while their benefit levels respond in opposite directions under tax and benefit adjust programs.

It appears that the pension systems we consider are effective to some degree in spreading shocks across generations, at least in some cases, although the effectiveness depends on the type of shock. For our base case, systems that incorporate adjustments to both taxes and benefits appear more robust to dealing with different types of shocks and to provide better performance overall, a result consistent with the findings in Auerbach and Lee (2011). These conclusions are tentative, given our specific modeling assumptions, but it seems clear from our results that the types of shocks hitting the economy should influence the nature of adjustments incorporated in a social security pension system. So, too, should preferences and the extent to which shocks affect factor prices, with tax adjustments appearing relatively less attractive in cases where the shocks increase the working population and depress wages.

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. Also, the use of this extra tool, at least for our version of the Swedish system, does not appear particularly effective in smoothing economic and demographic shocks under our base case assumptions, although it performs much better in alternative scenarios. These findings suggest that further research is needed regarding the design of unfunded systems that allow short-run fluctuations in annual pay-as-you-go balance.

References

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.

Gourinchas, Pierre-Olivier, and Jonathan Parker (2002) "Consumption over the Life Cycle," *Econometrica*, 70(1), January, 47-89.

Lee, Ronald and Shripad Tuljapurkar (1994) "Stochastic Population Projections for the United States: Beyond High, Medium and Low," *Journal of the American Statistical Association*, 89(428), December, 1175-1189.

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)

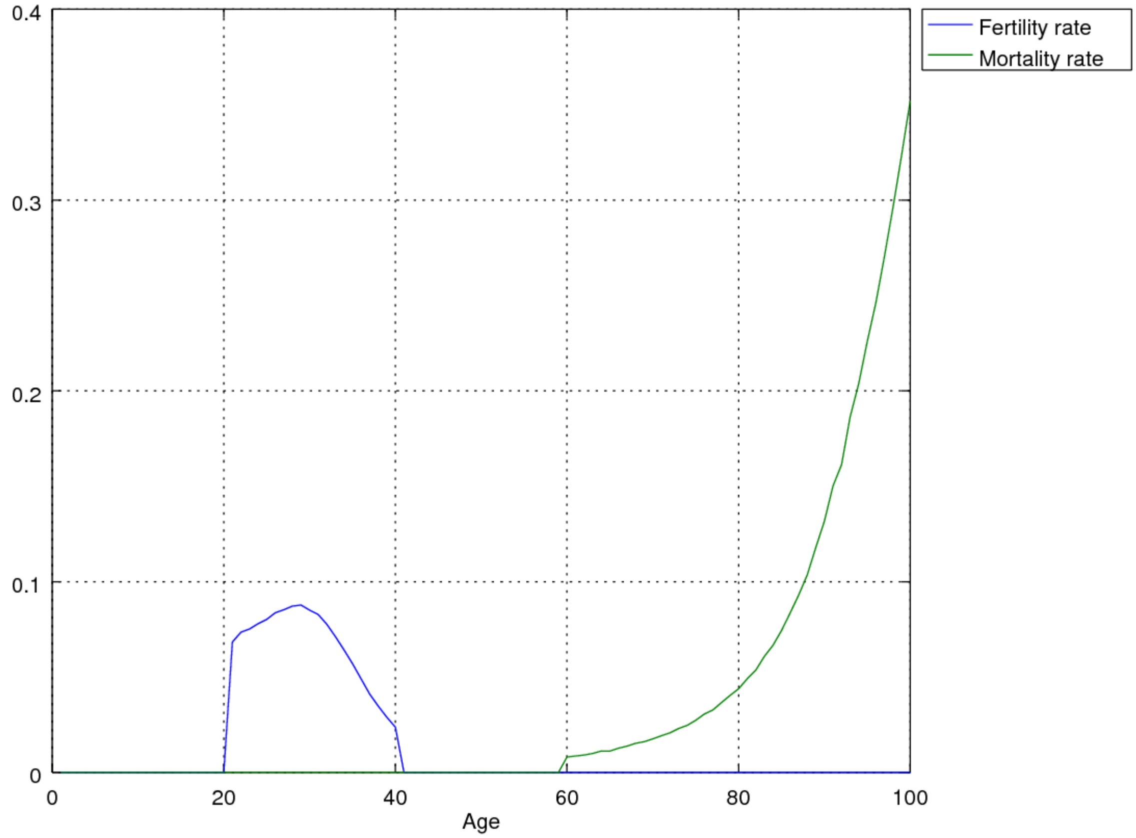

Figure 2. Population by Age, Fertility Shock
(adjusted for trend population growth)

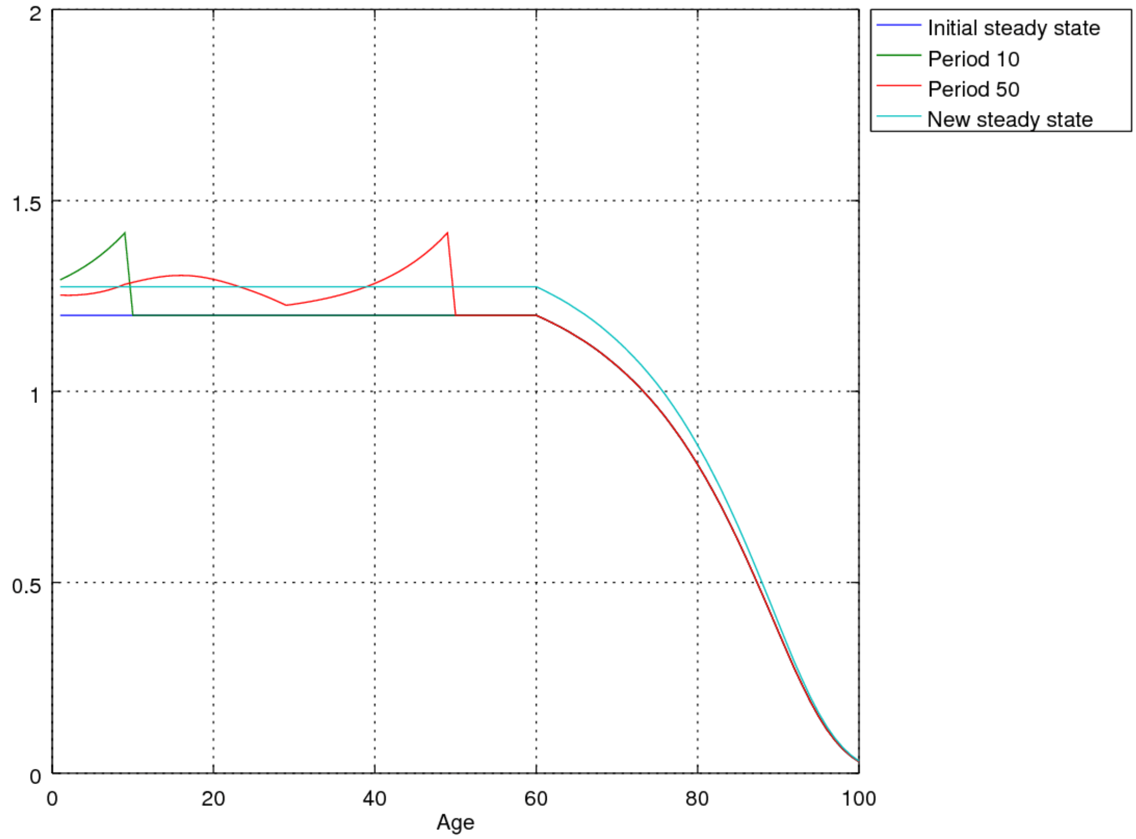

Figure 3. Aggregate Population (Normalized by Initial Population) Over Time, Fertility Shock

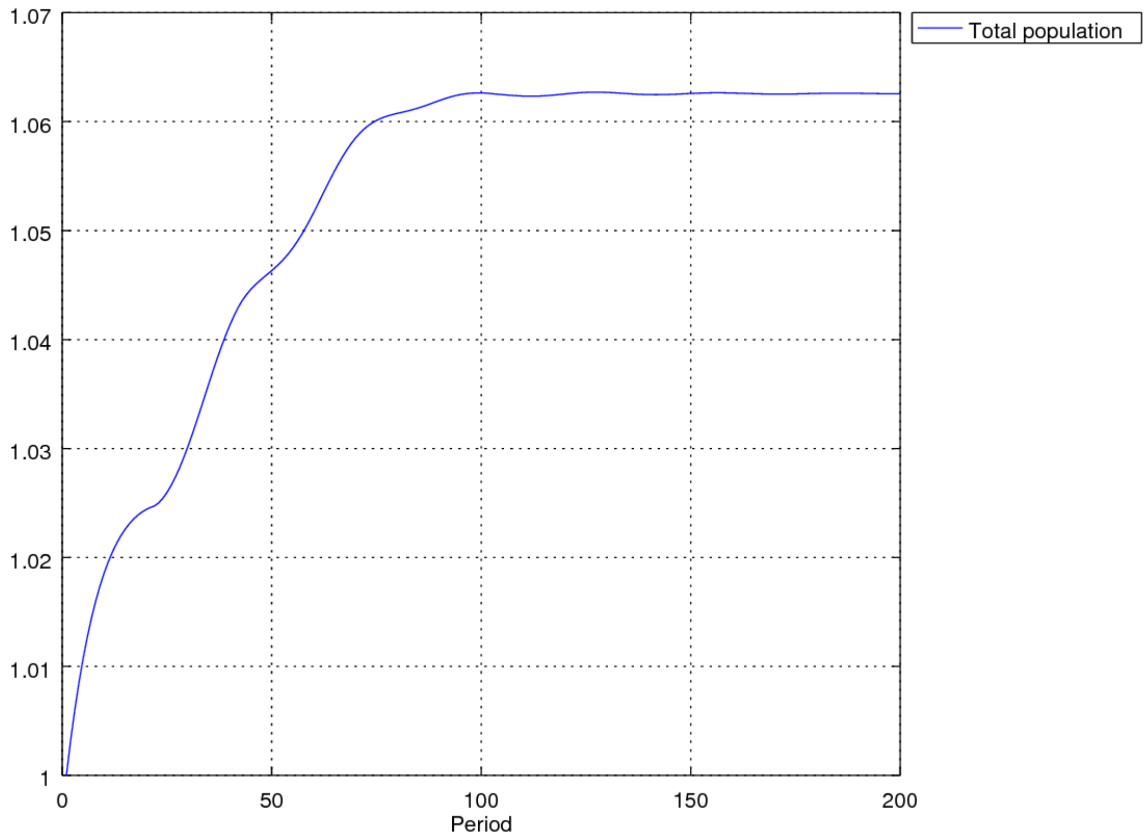

Figure 4. Population by Age, Mortality Shock
(adjusted for trend population growth)

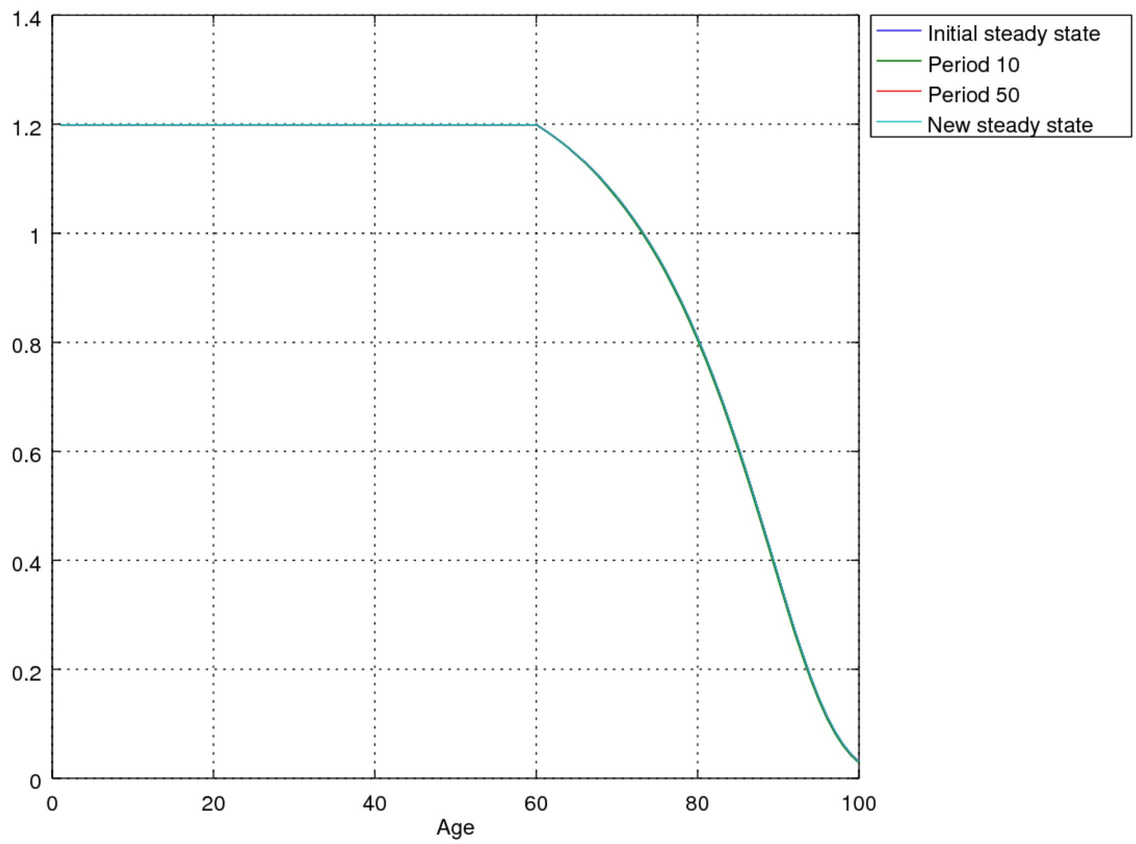

Figure 5. Aggregate Population (Normalized by Initial Population) Over Time, Mortality Shock

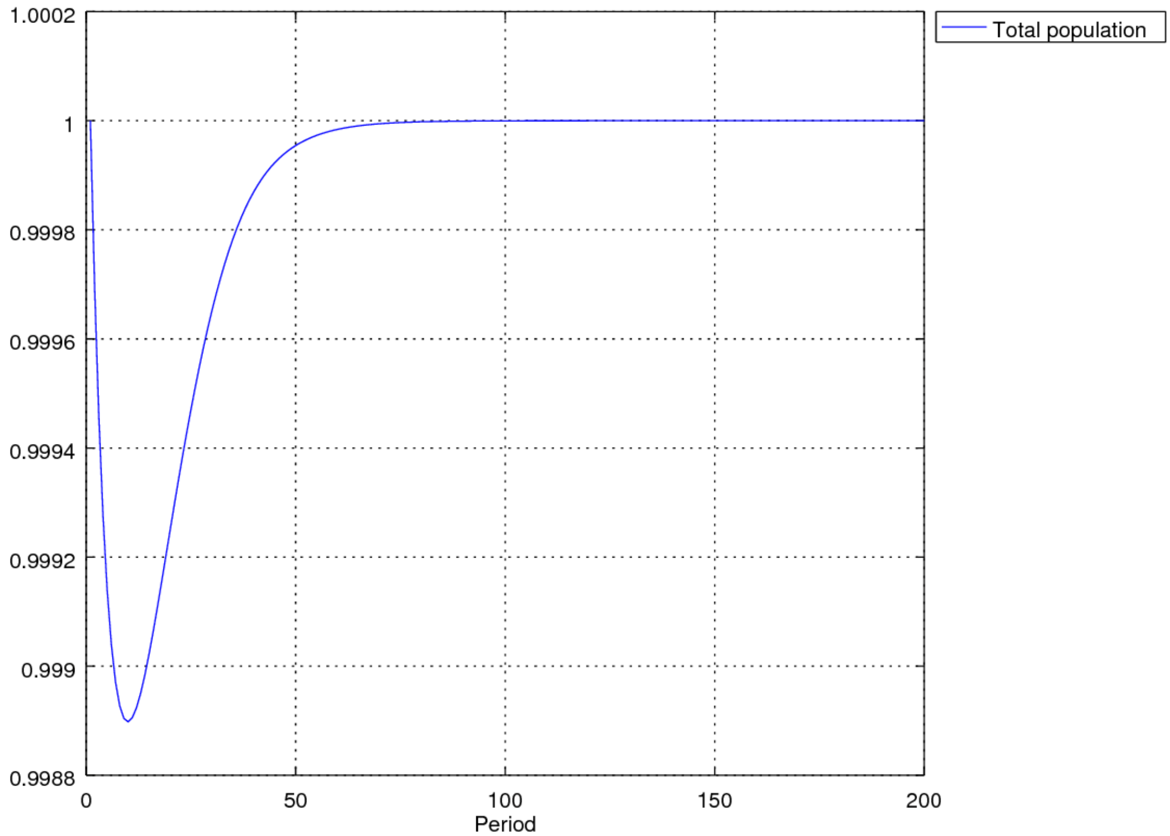

Figure 6a. Productivity Shock: Wealth Equivalents by Year of Adulthood

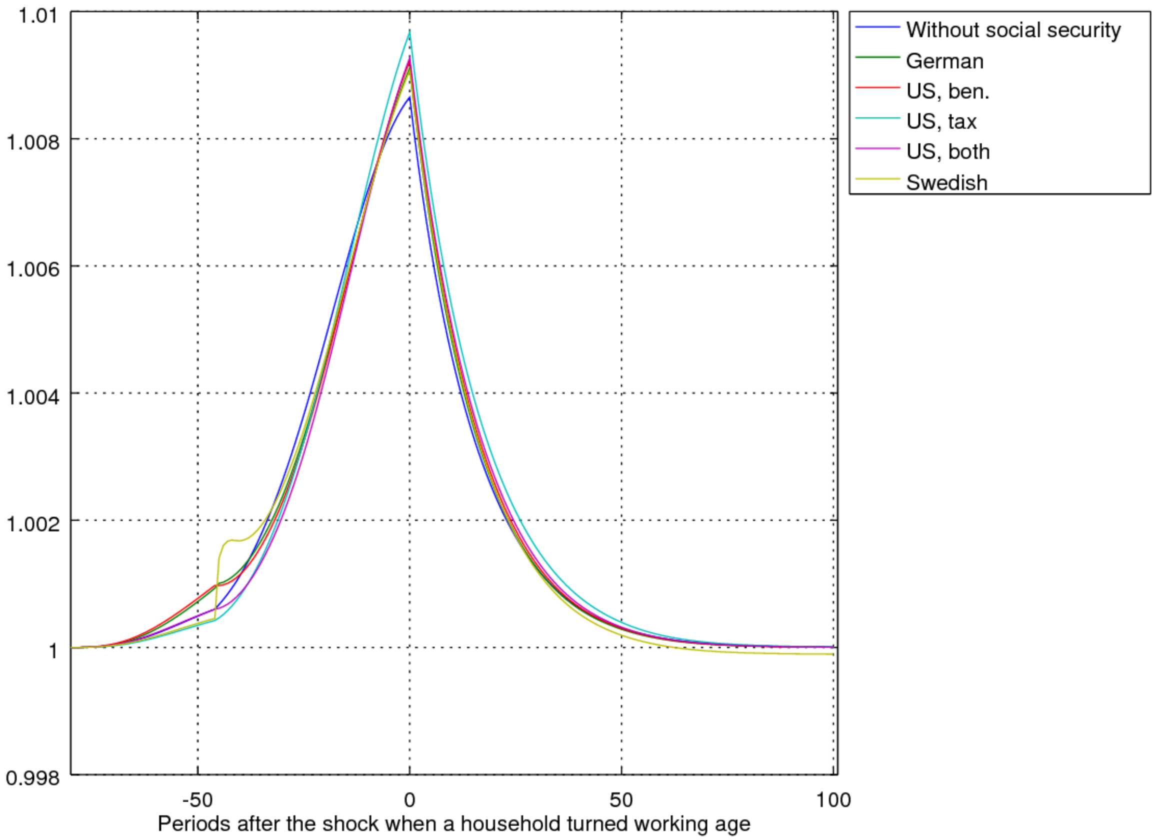

Wealth Equivalents: Variance (Relative to Without Social Security)

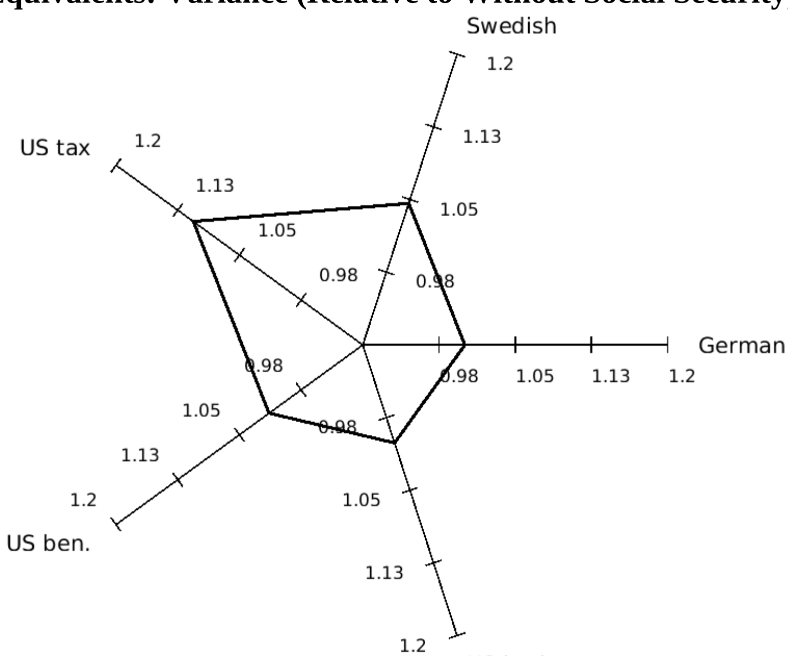

Figure 6b. Productivity Shock: Aggregate Wage Rate by Year

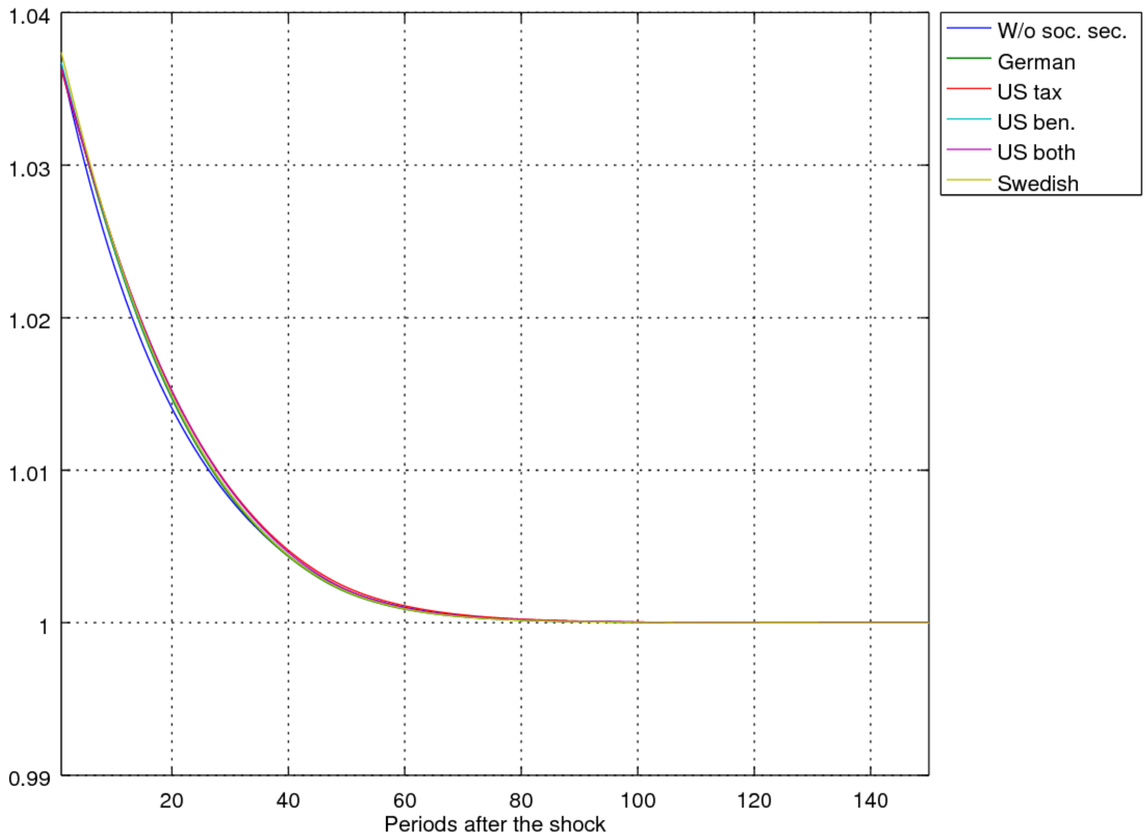

Figure 6c. Productivity Shock: Interest Rate by Year

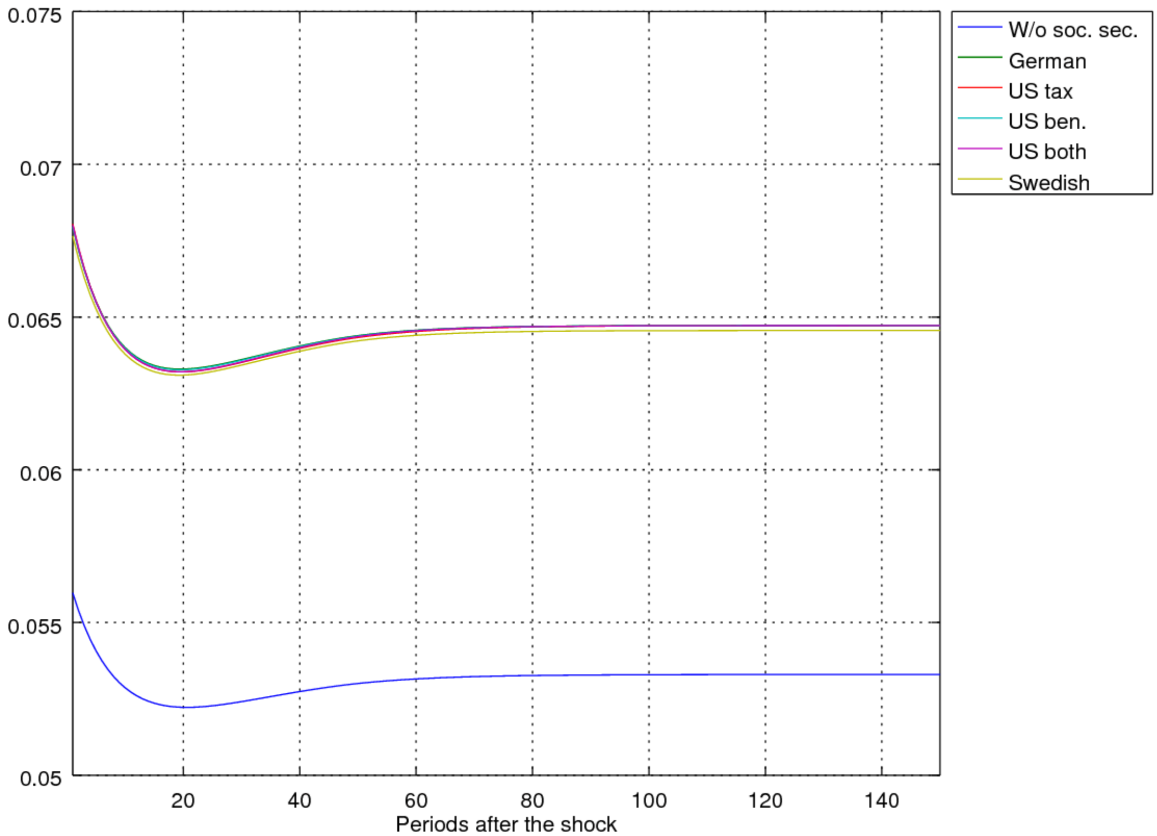

Figure 6d. Productivity Shock: Social Security Tax Rate by Year

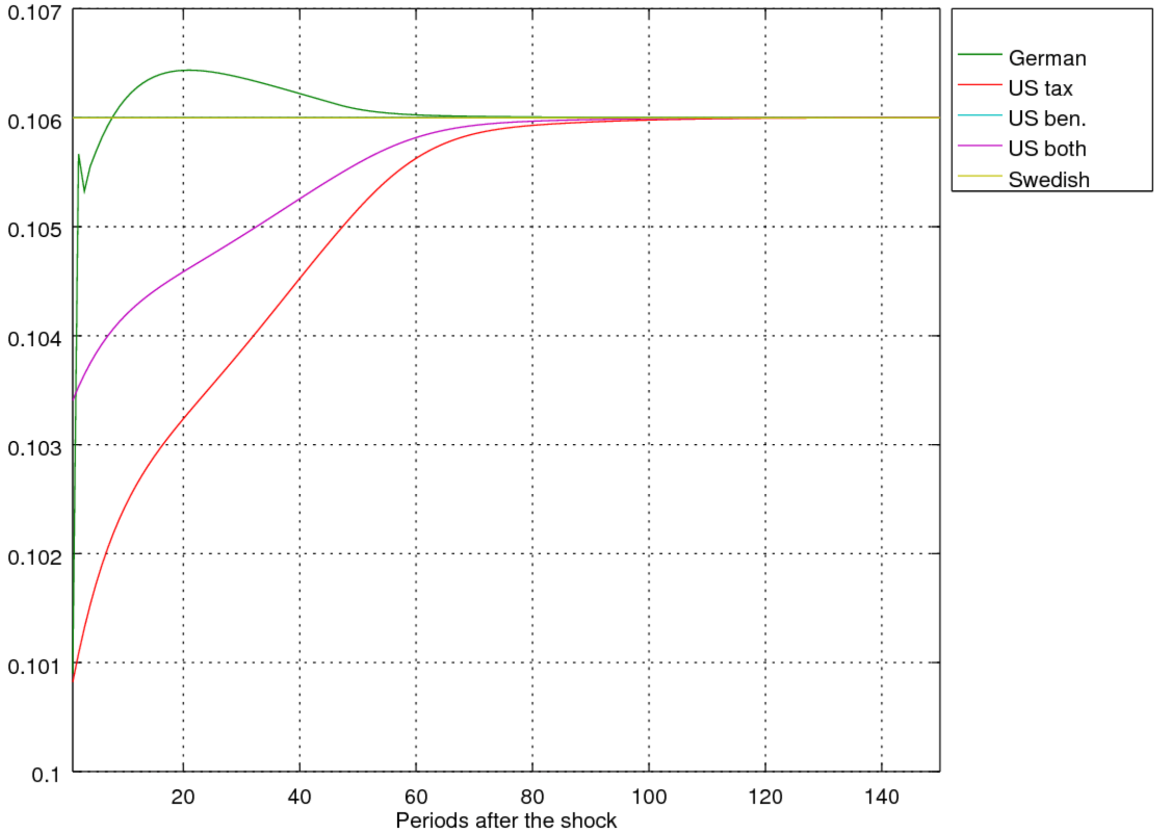

Figure 7a. Mortality Shock: Wealth Equivalents by Year of Adulthood

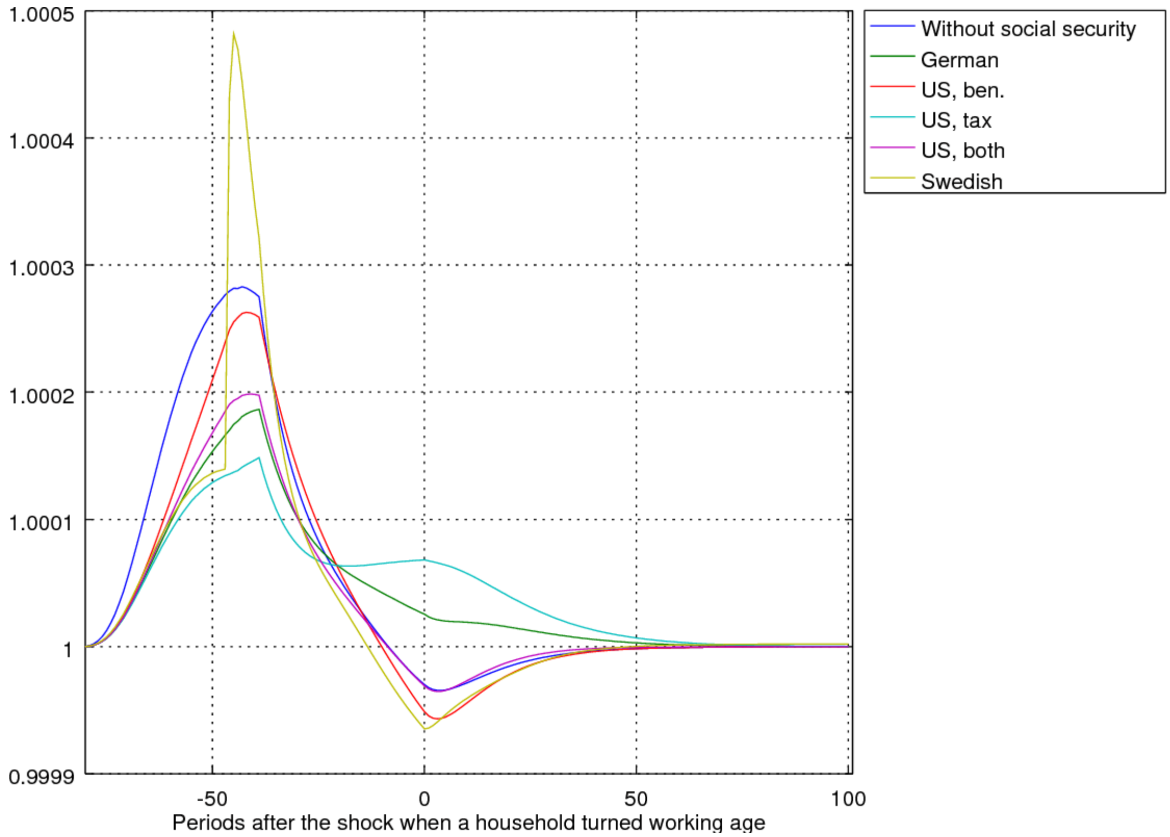

Wealth Equivalents: Variance (Relative to Without Social Security)

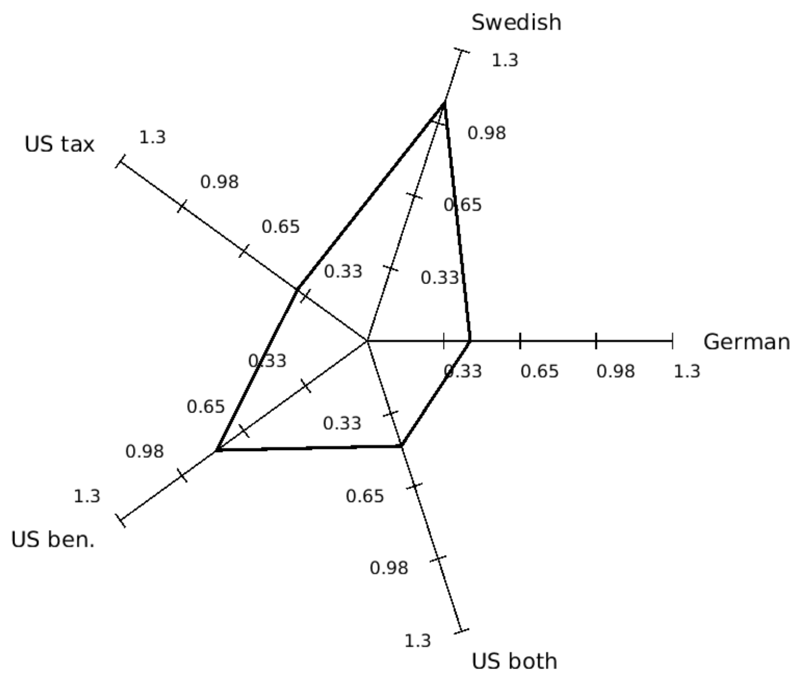

Figure 7b. Mortality Shock: Aggregate Wage Rate by Year

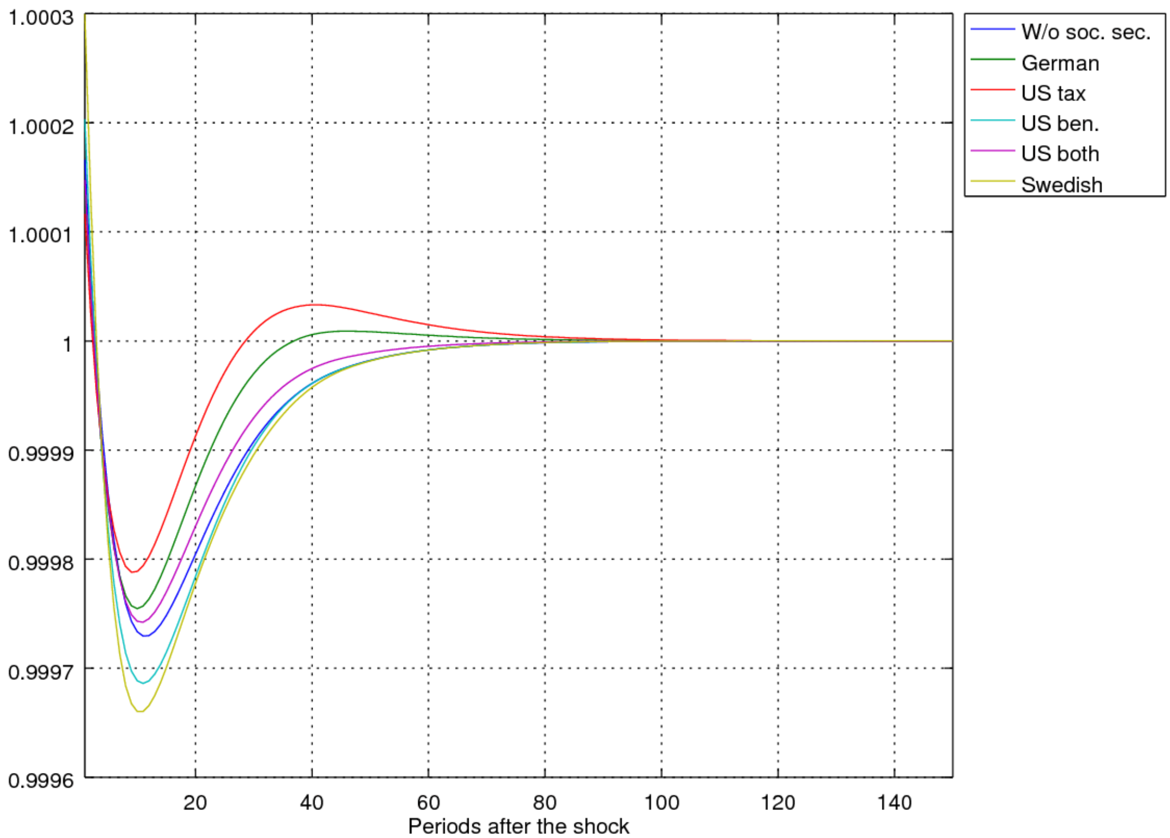

Figure 7c. Mortality Shock: Interest Rate by Year

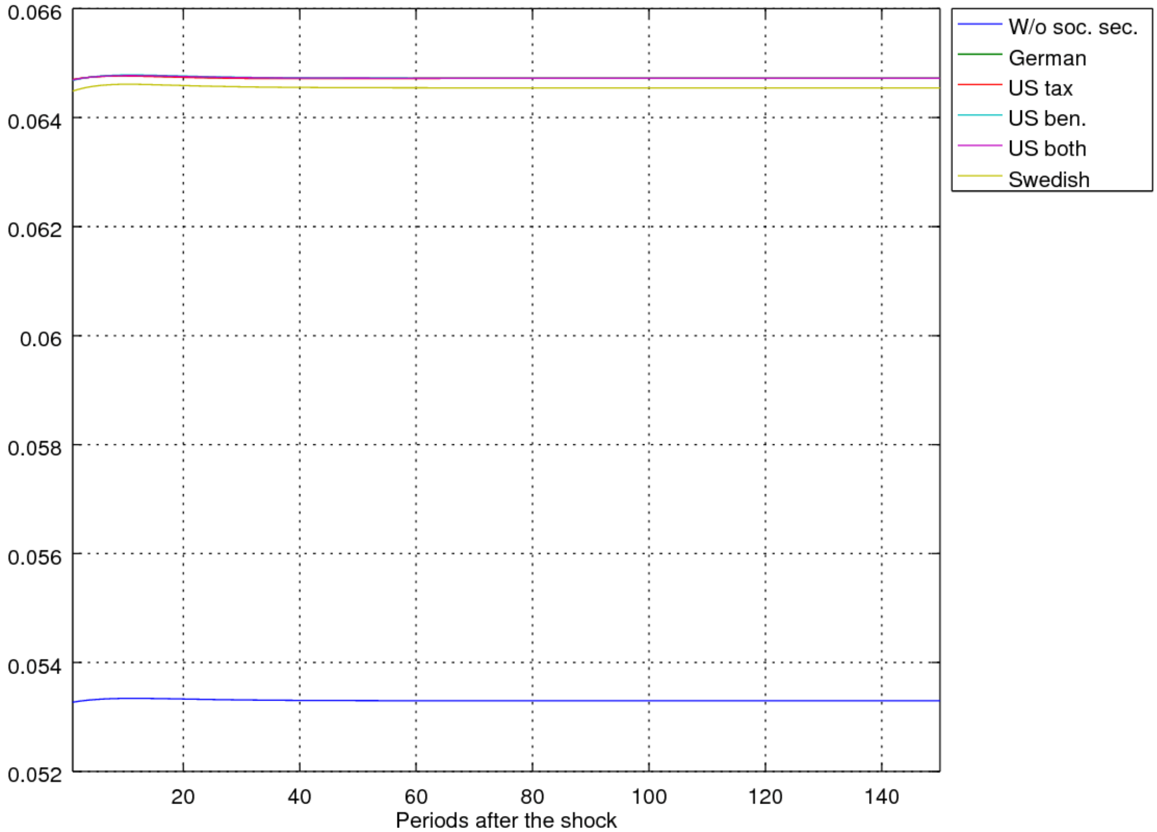

Figure 7d. Mortality Shock: Social Security Tax Rate by Year

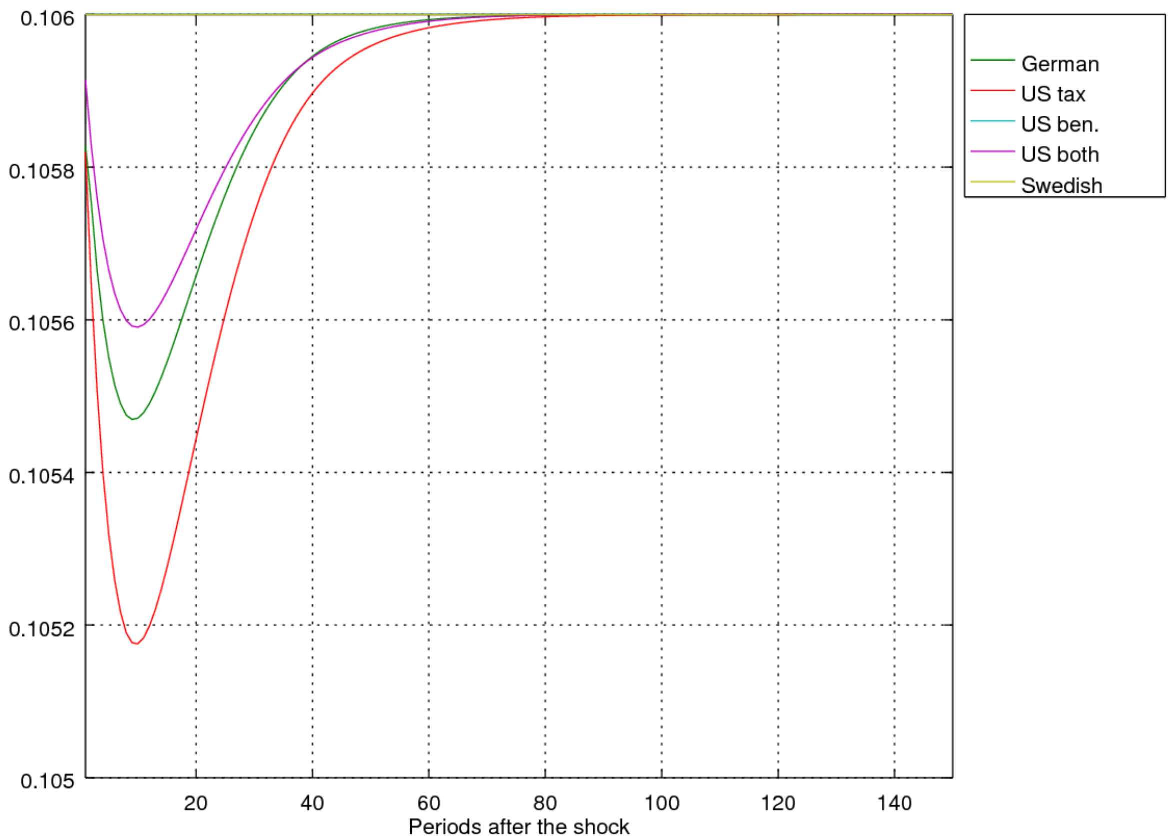

Figure 8a. Fertility Shock: Wealth Equivalents by Year of Adulthood

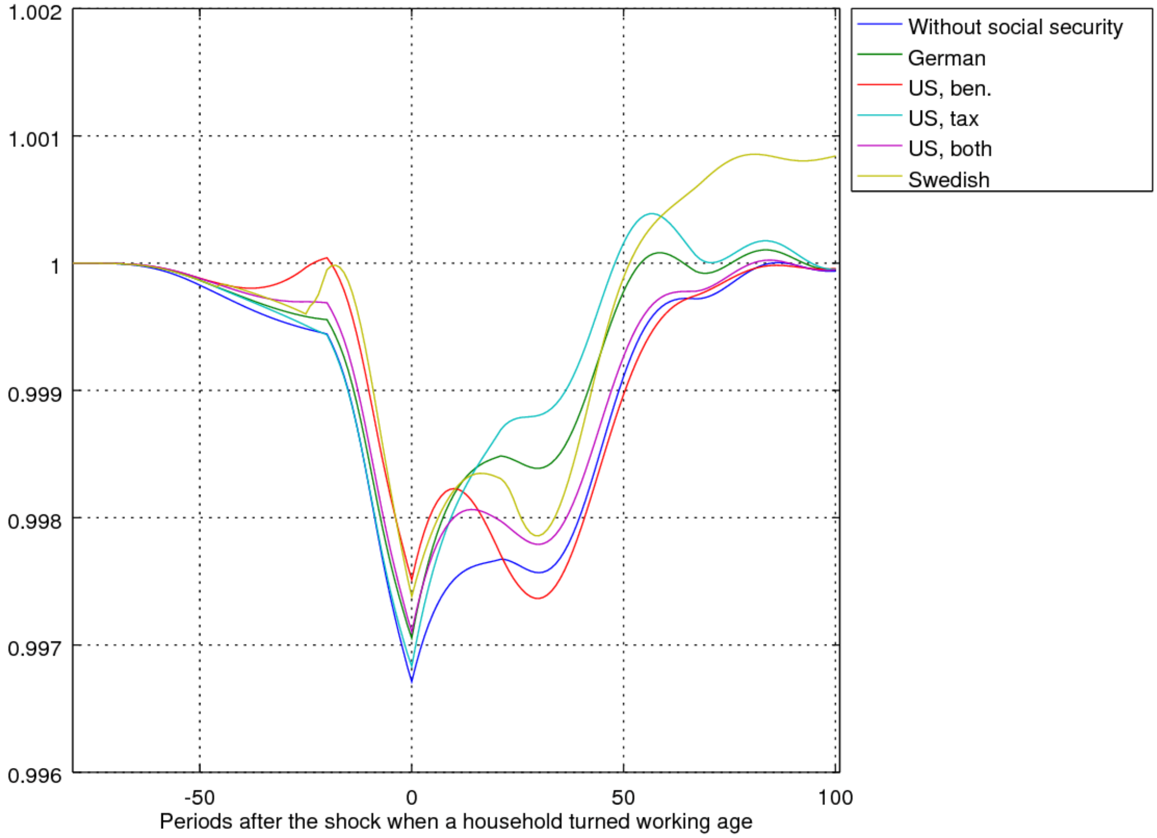

Wealth Equivalents: Variance (Relative to Without Social Security)

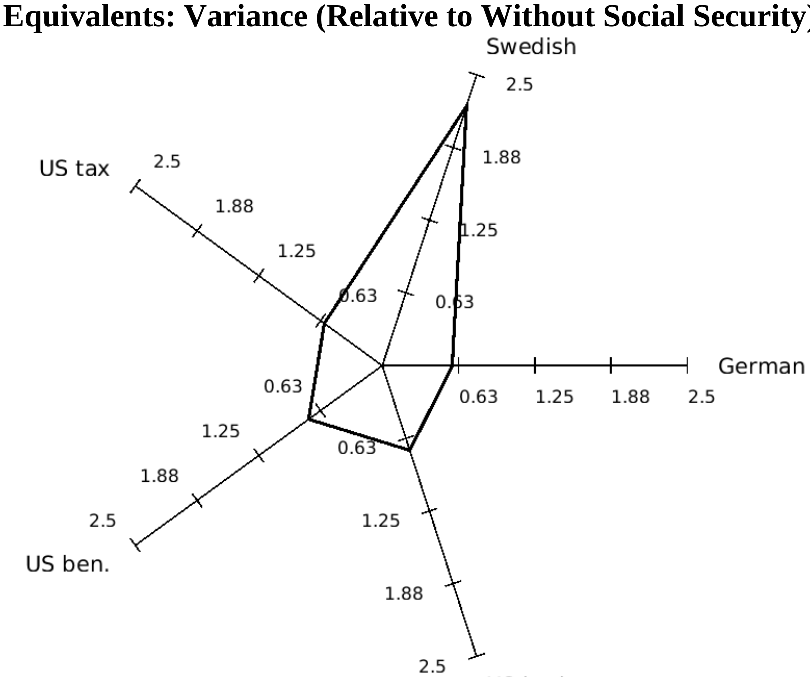

Figure 8b. Fertility Shock: Aggregate Wage Rate by Year

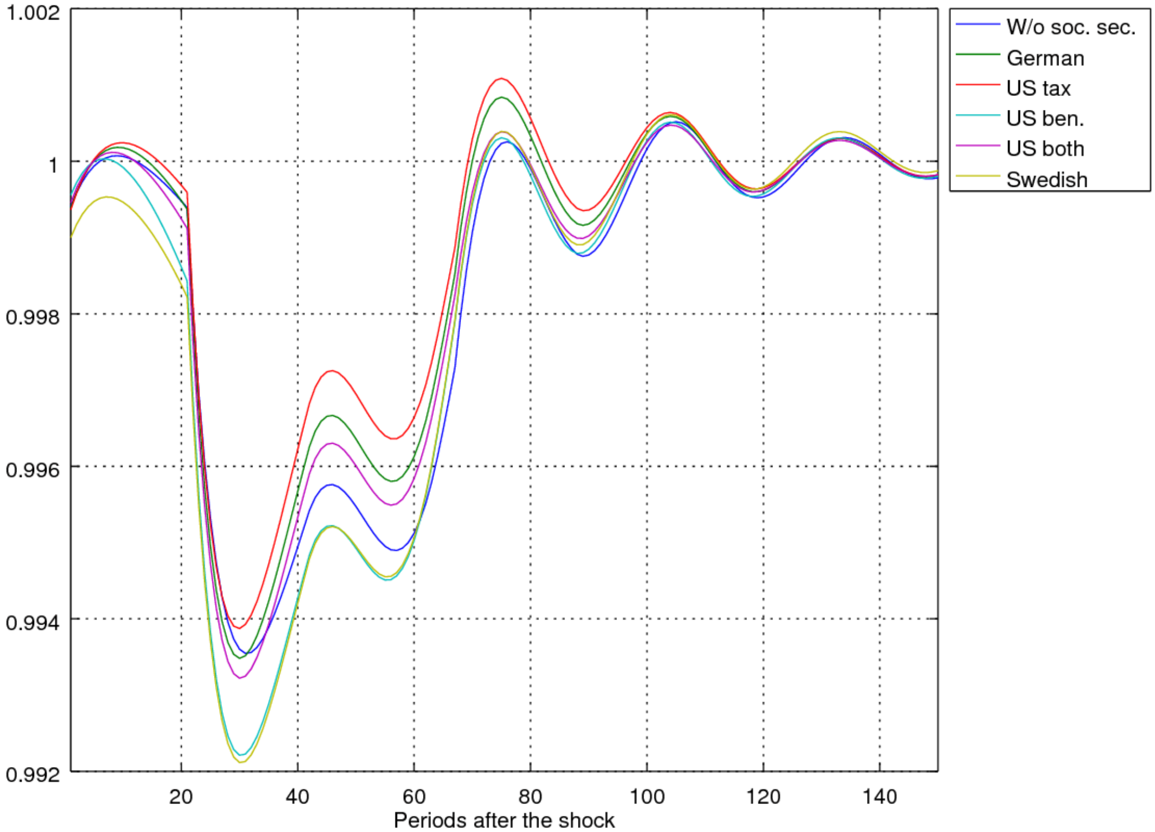

Figure 8c. Fertility Shock: Interest Rate by Year

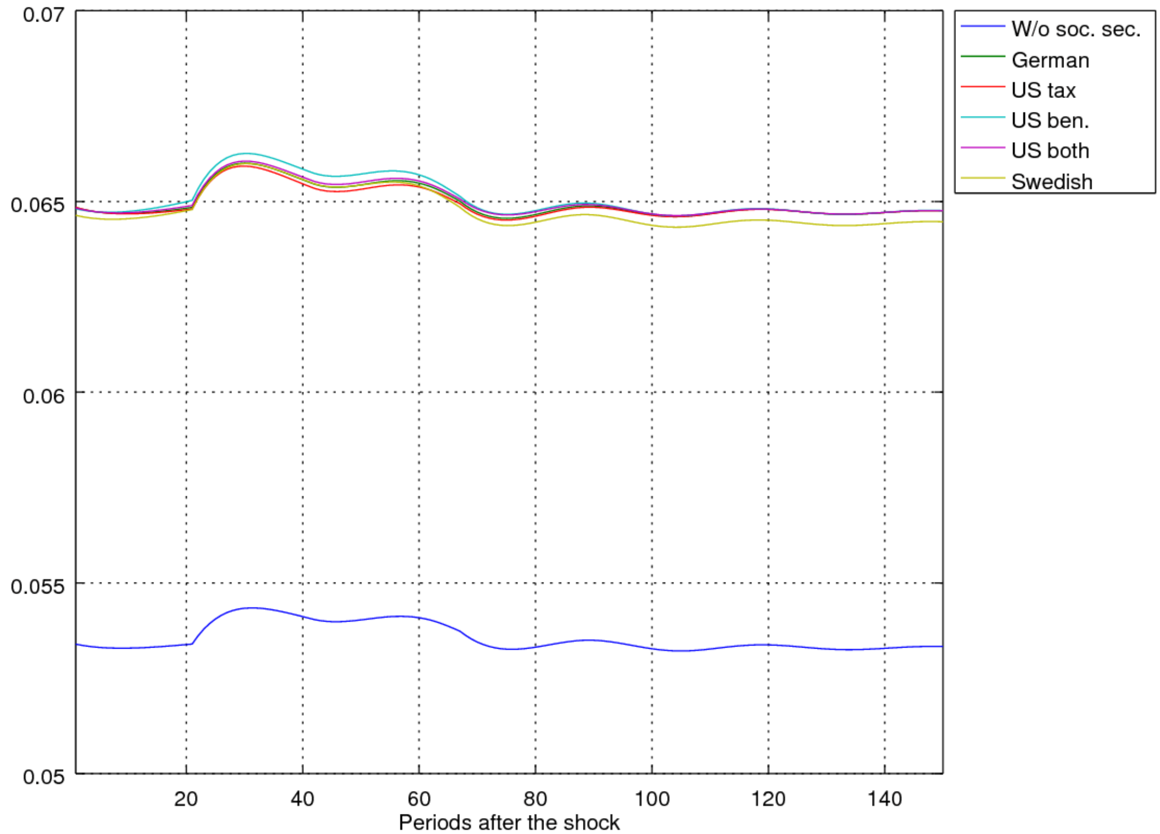

Figure 8d. Fertility Shock: Social Security Tax Rate by Year

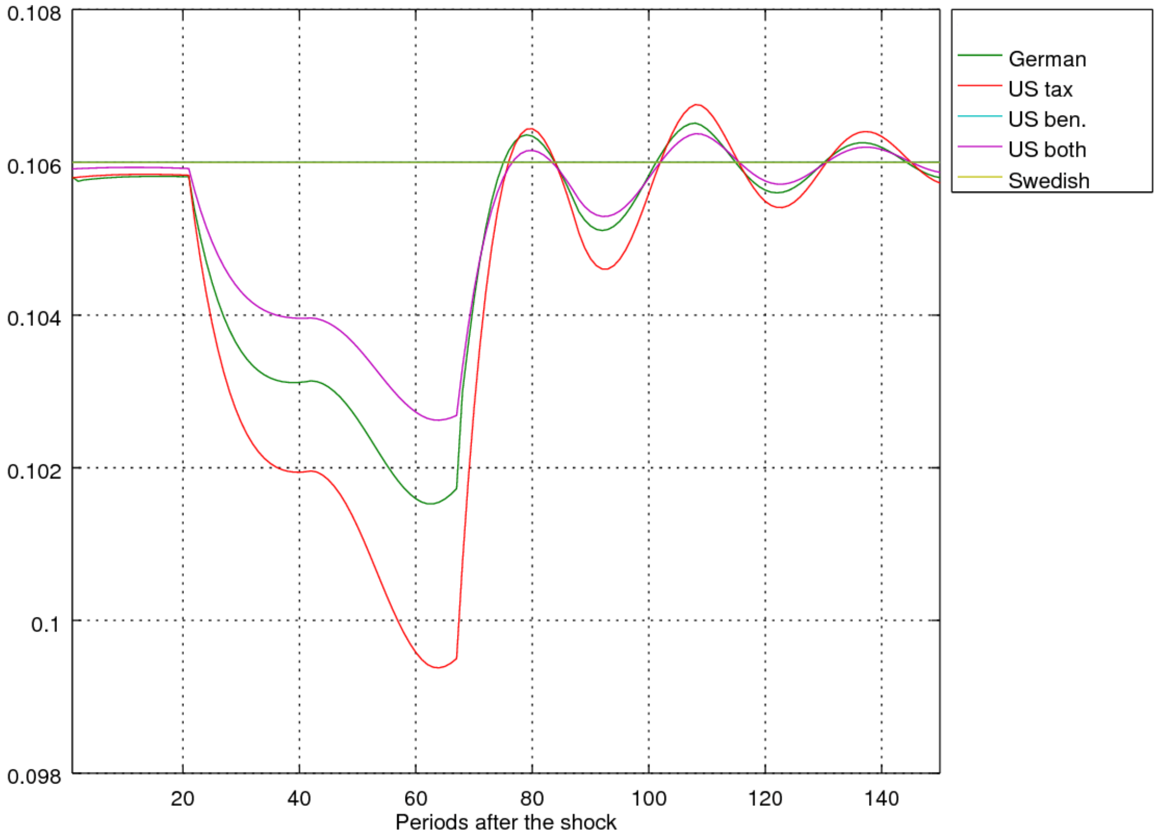

Figure 9. Migration: Age Profile of Immigrant Flow (to Germany in 2014)

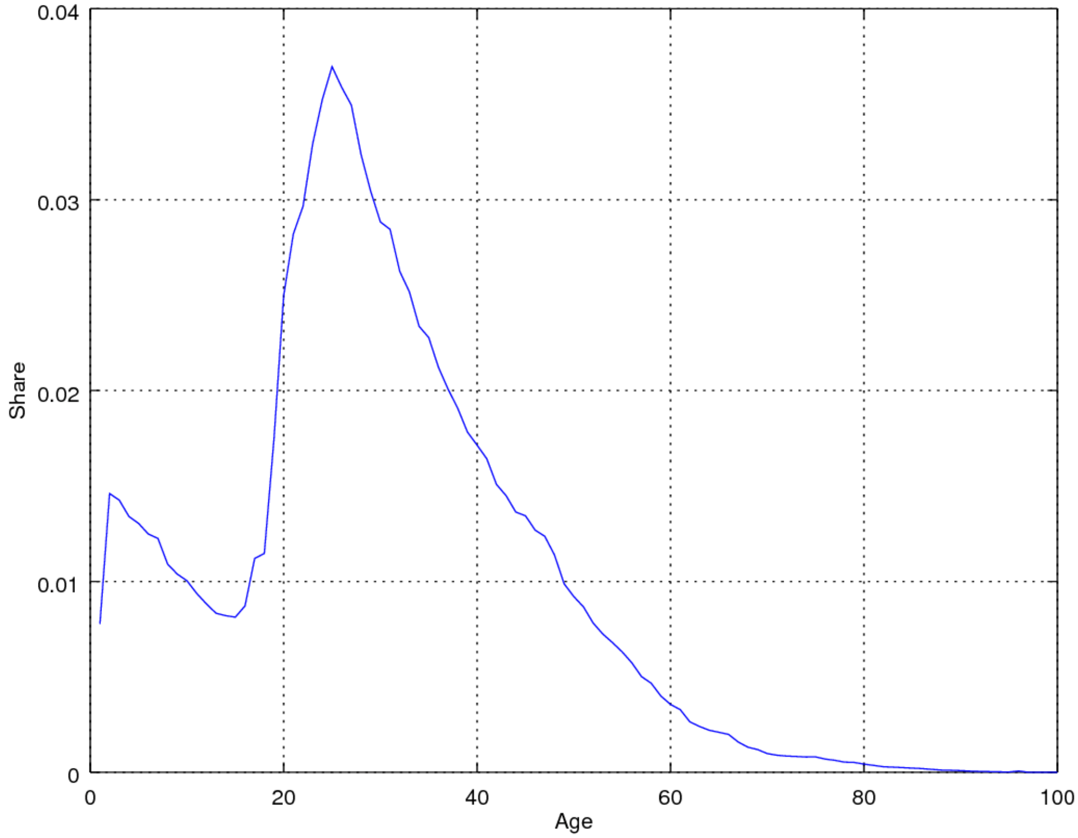

Figure 10. Population by Age, Migration Shock
(adjusted for trend population growth)

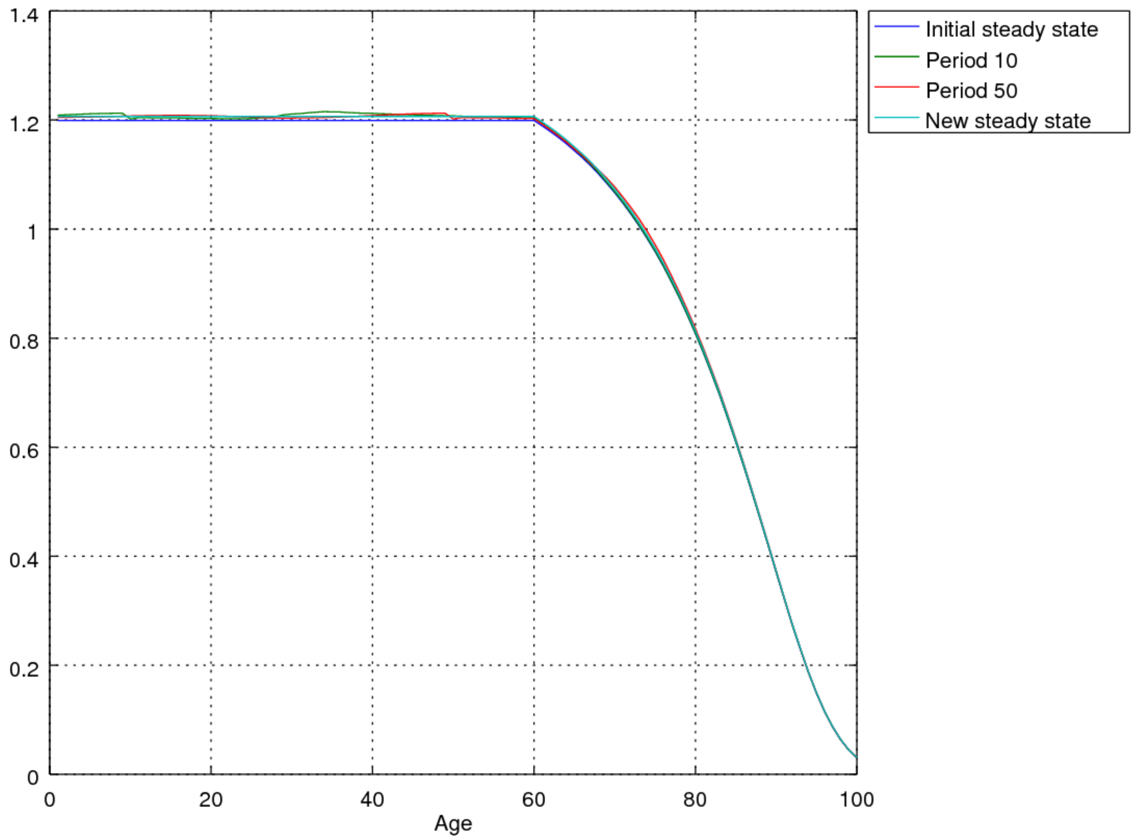

Figure 11. Aggregate Population (Normalized by Initial Population) Over Time, Migration Shock

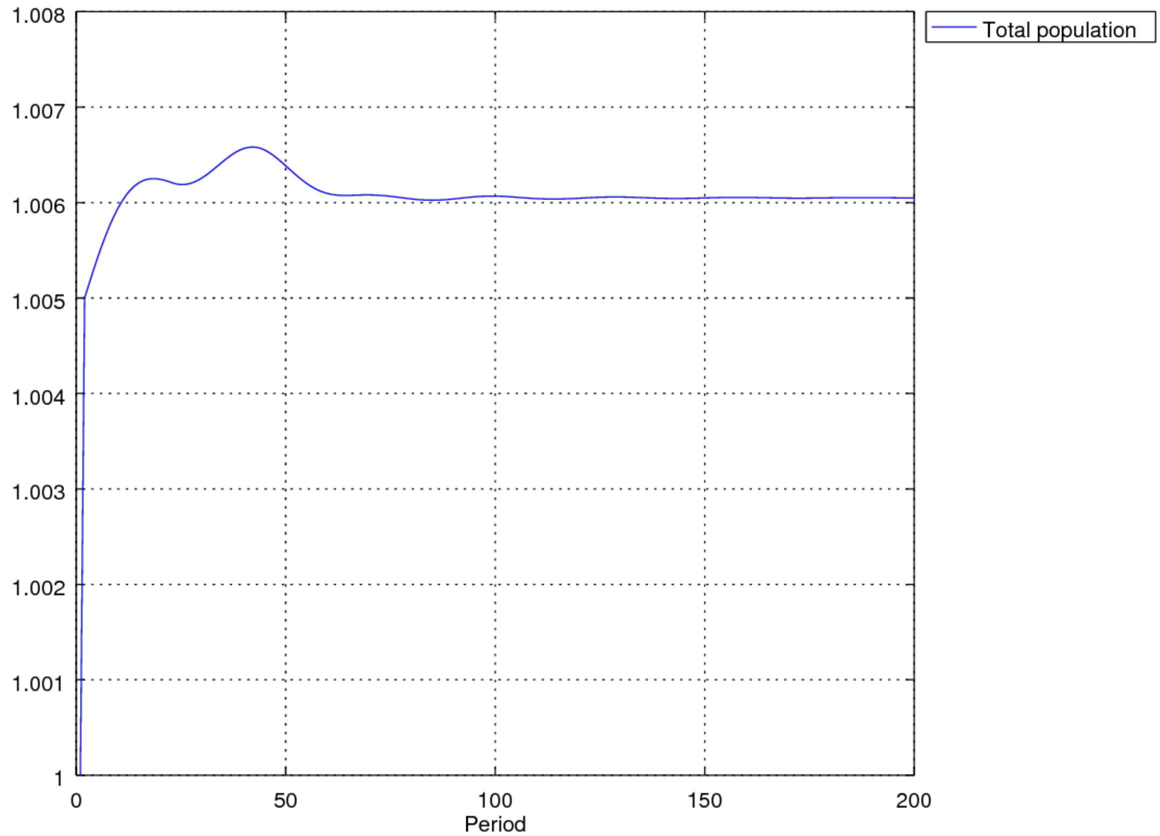

Figure 12a. Migration Shock: Wealth Equivalents by Year of Adulthood

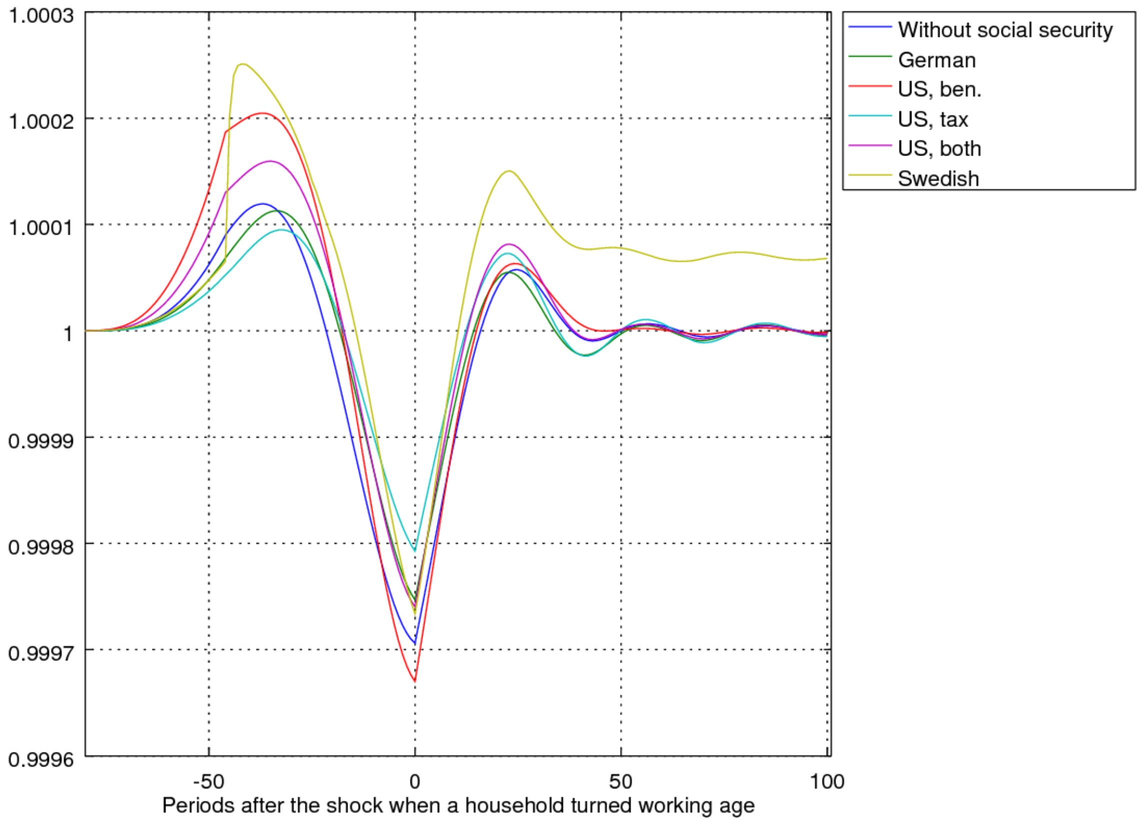

Wealth Equivalents: Variance (Relative to Without Social Security)

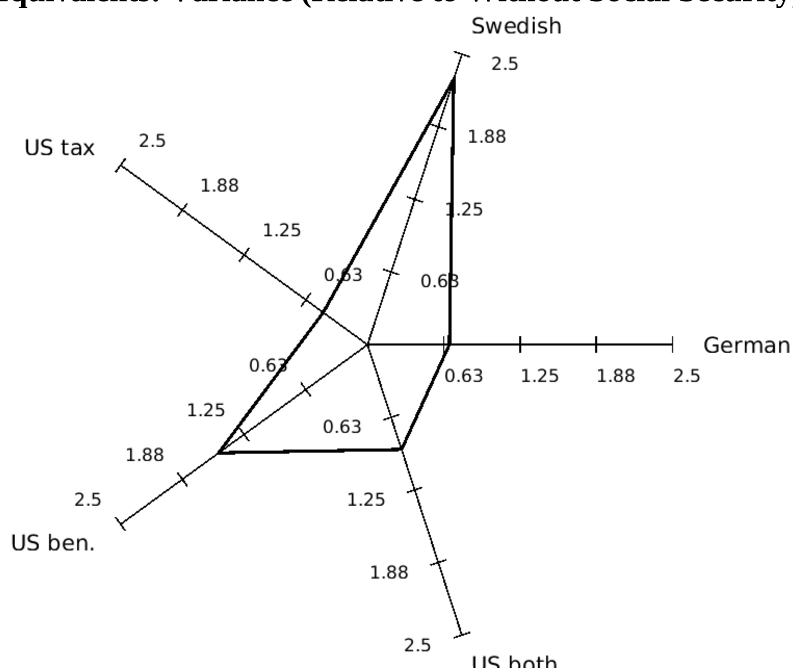

Figure 12b. Migration Shock: Aggregate Wage Rate by Year

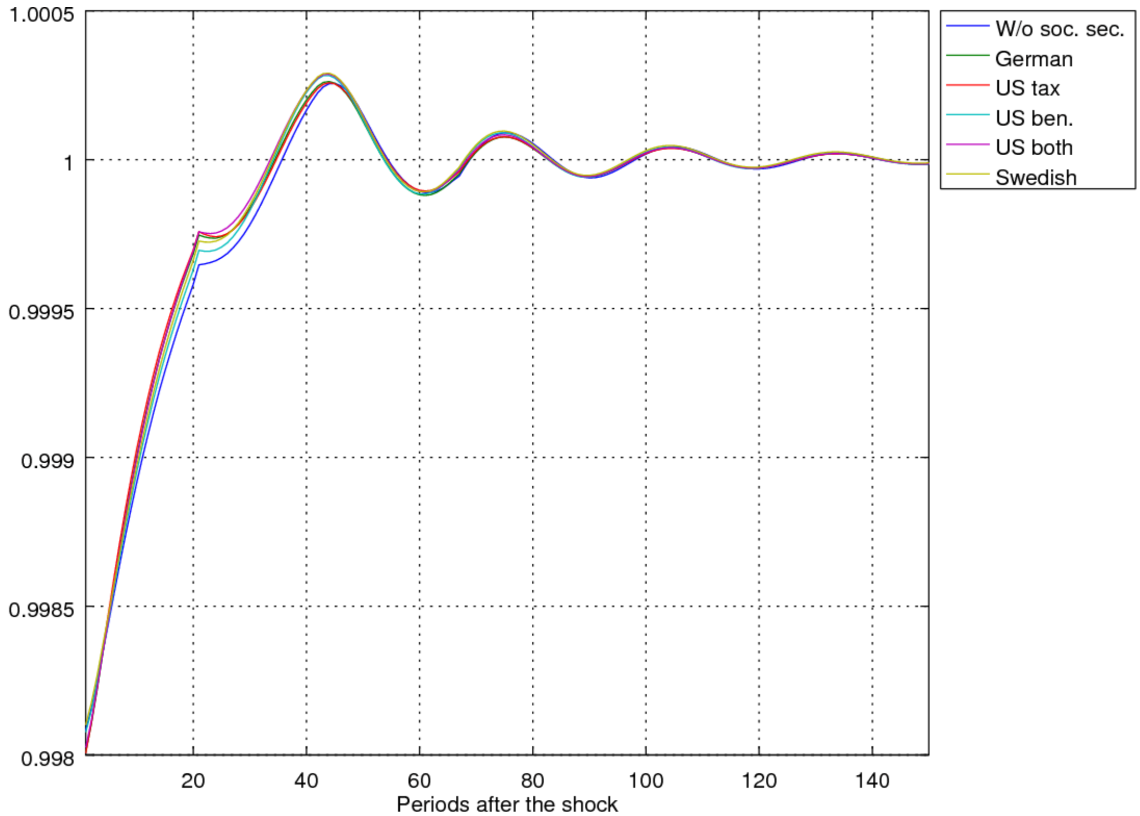

Figure 12c. Migration Shock: Interest Rate by Year

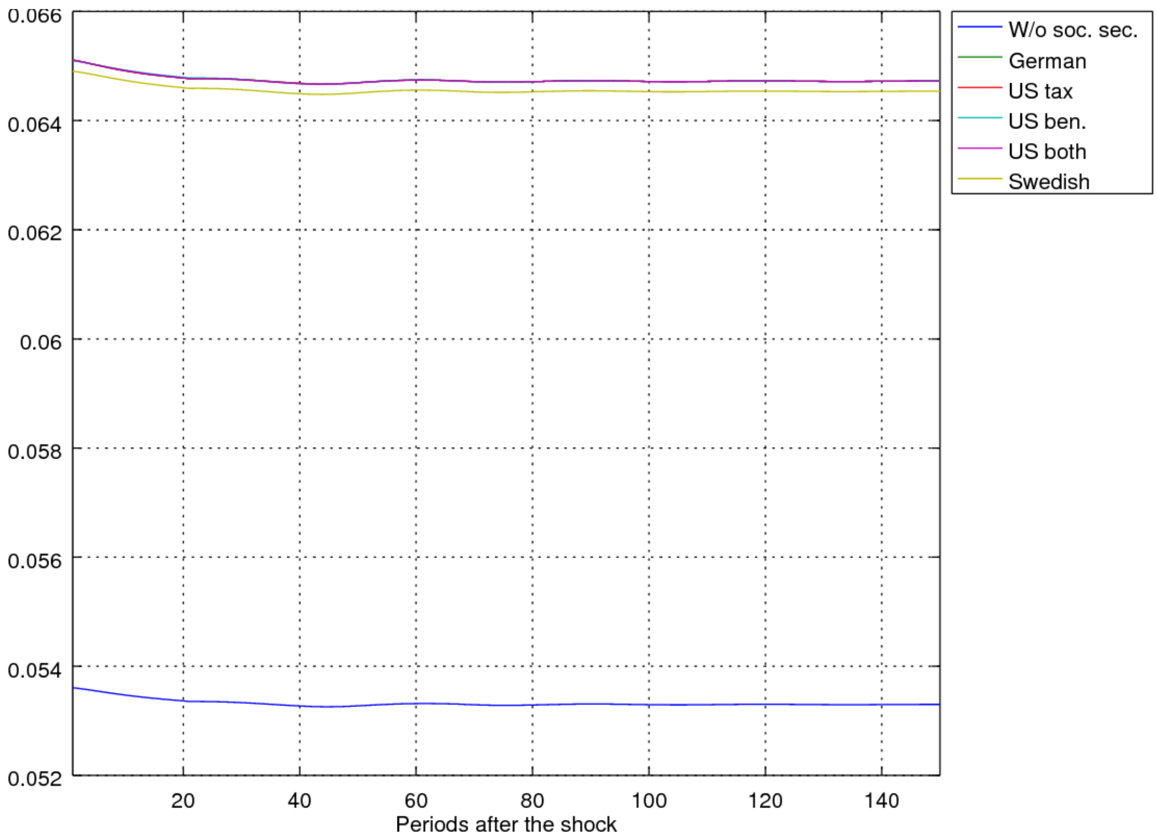

Figure 12d. Migration Shock: Social Security Tax Rate by Year

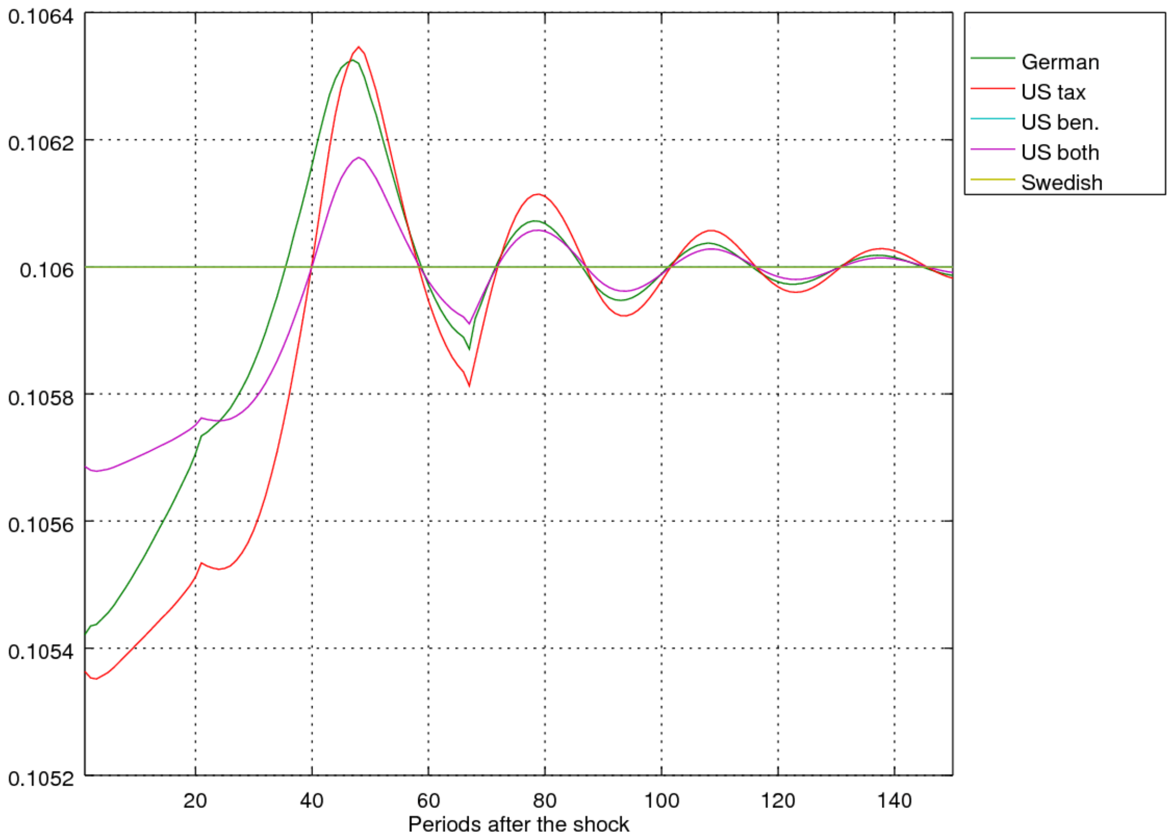

Figure 13. Balance Ratio, Swedish System

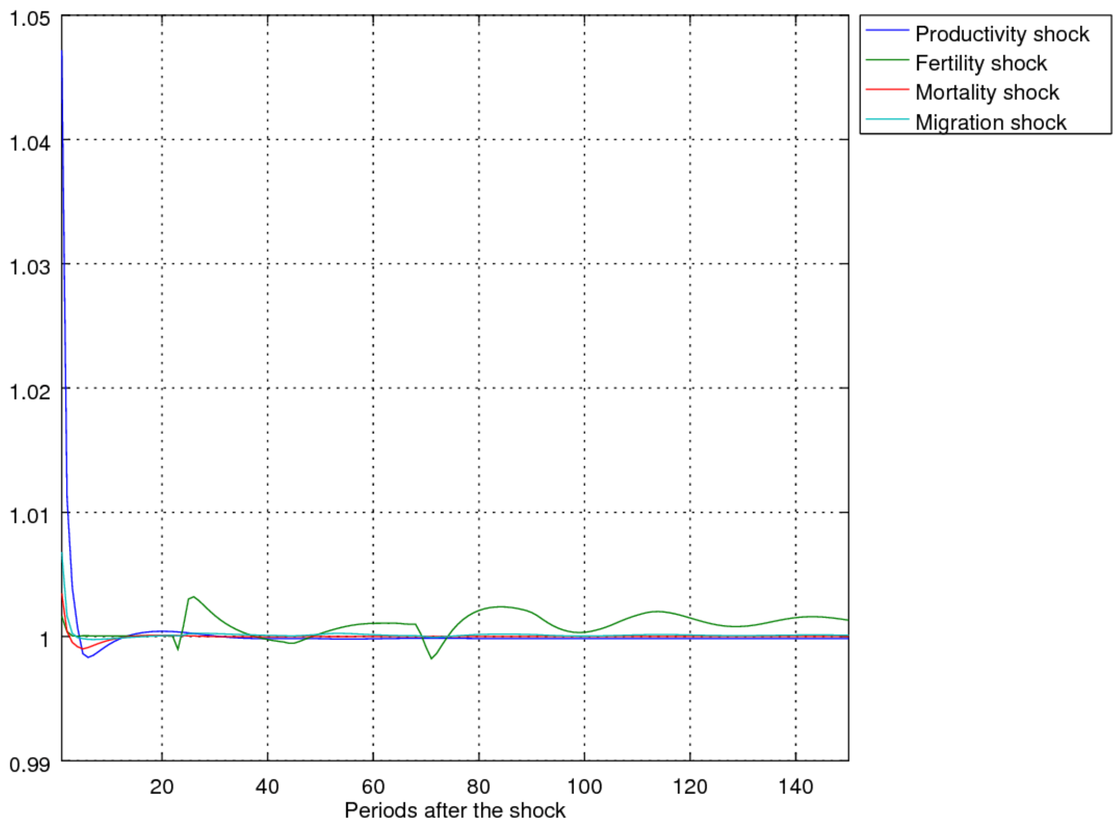

Figure 14. Fund-Wages Ratio, Swedish System

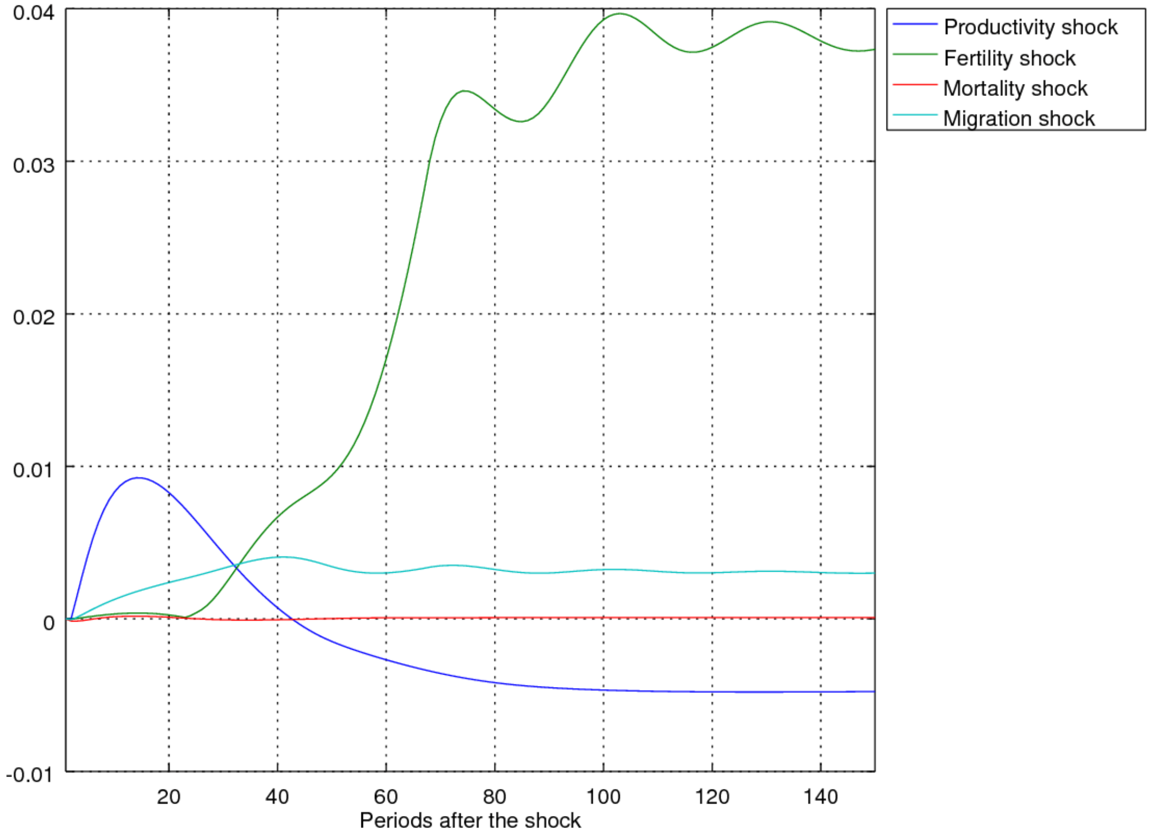

Figure 15. Productivity Shock: Wealth Equivalents for γ = 0.35 (upper), 0.7 (lower)

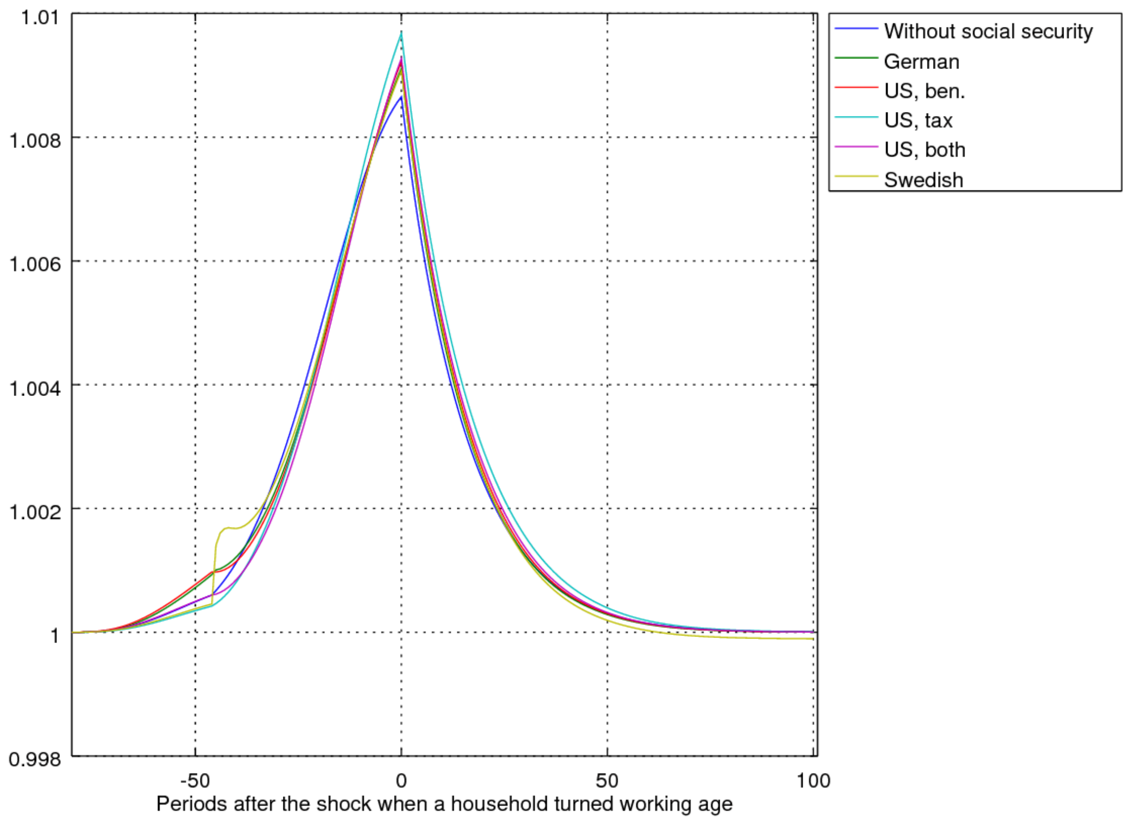

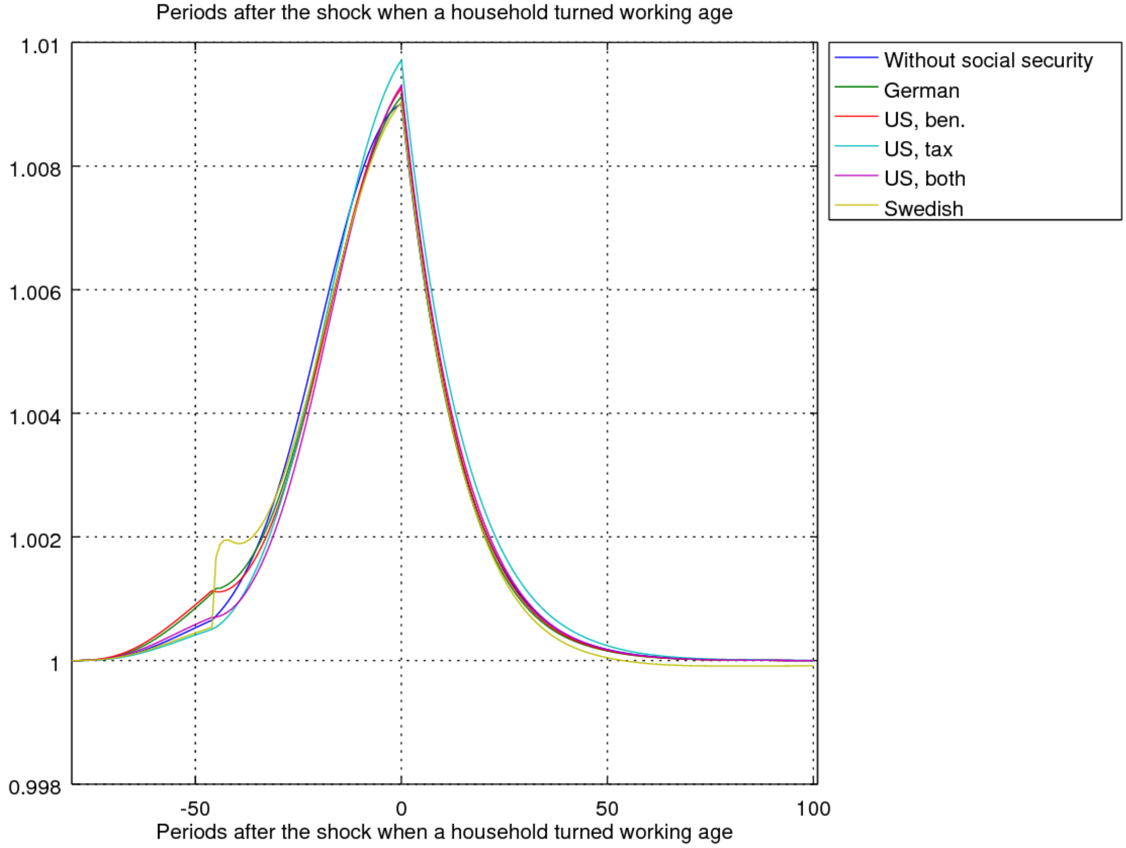

Figure 16. Mortality Shock: Wealth Equivalents for γ = 0.35 (upper), 0.7 (lower)

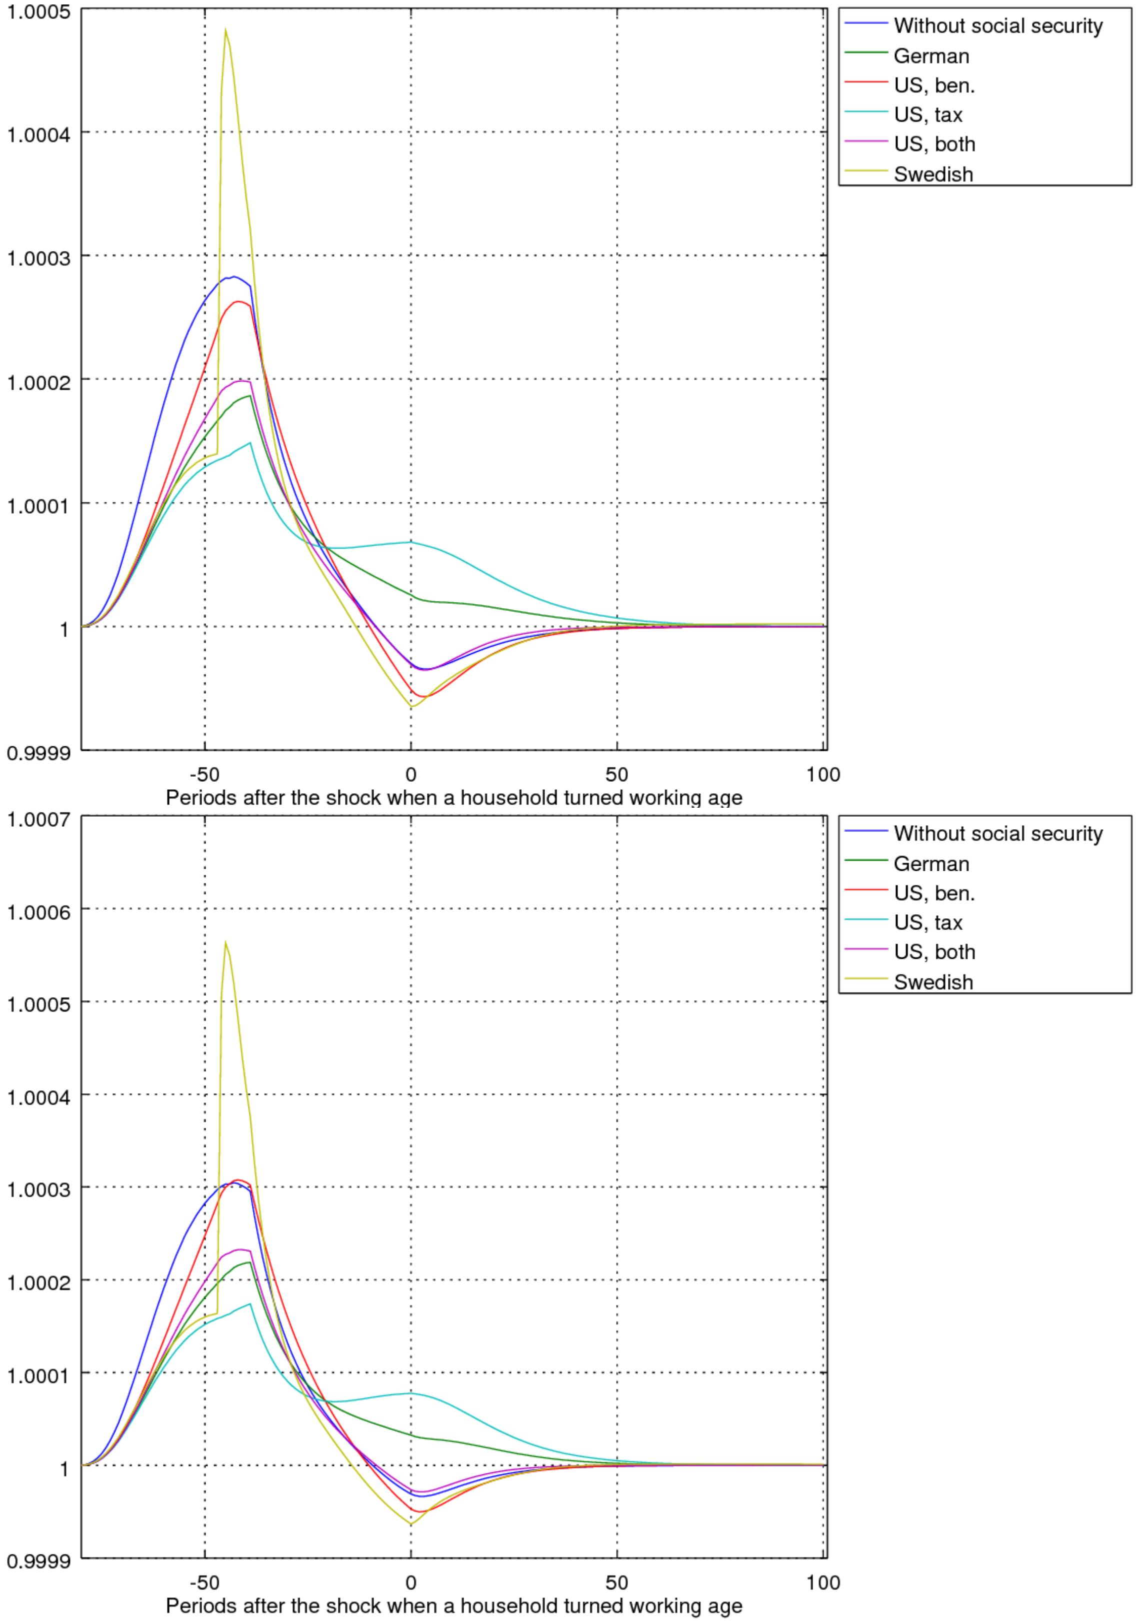

Figure 17. Fertility Shock: Wealth Equivalents for γ = 0.35 (upper), 0.7 (lower)

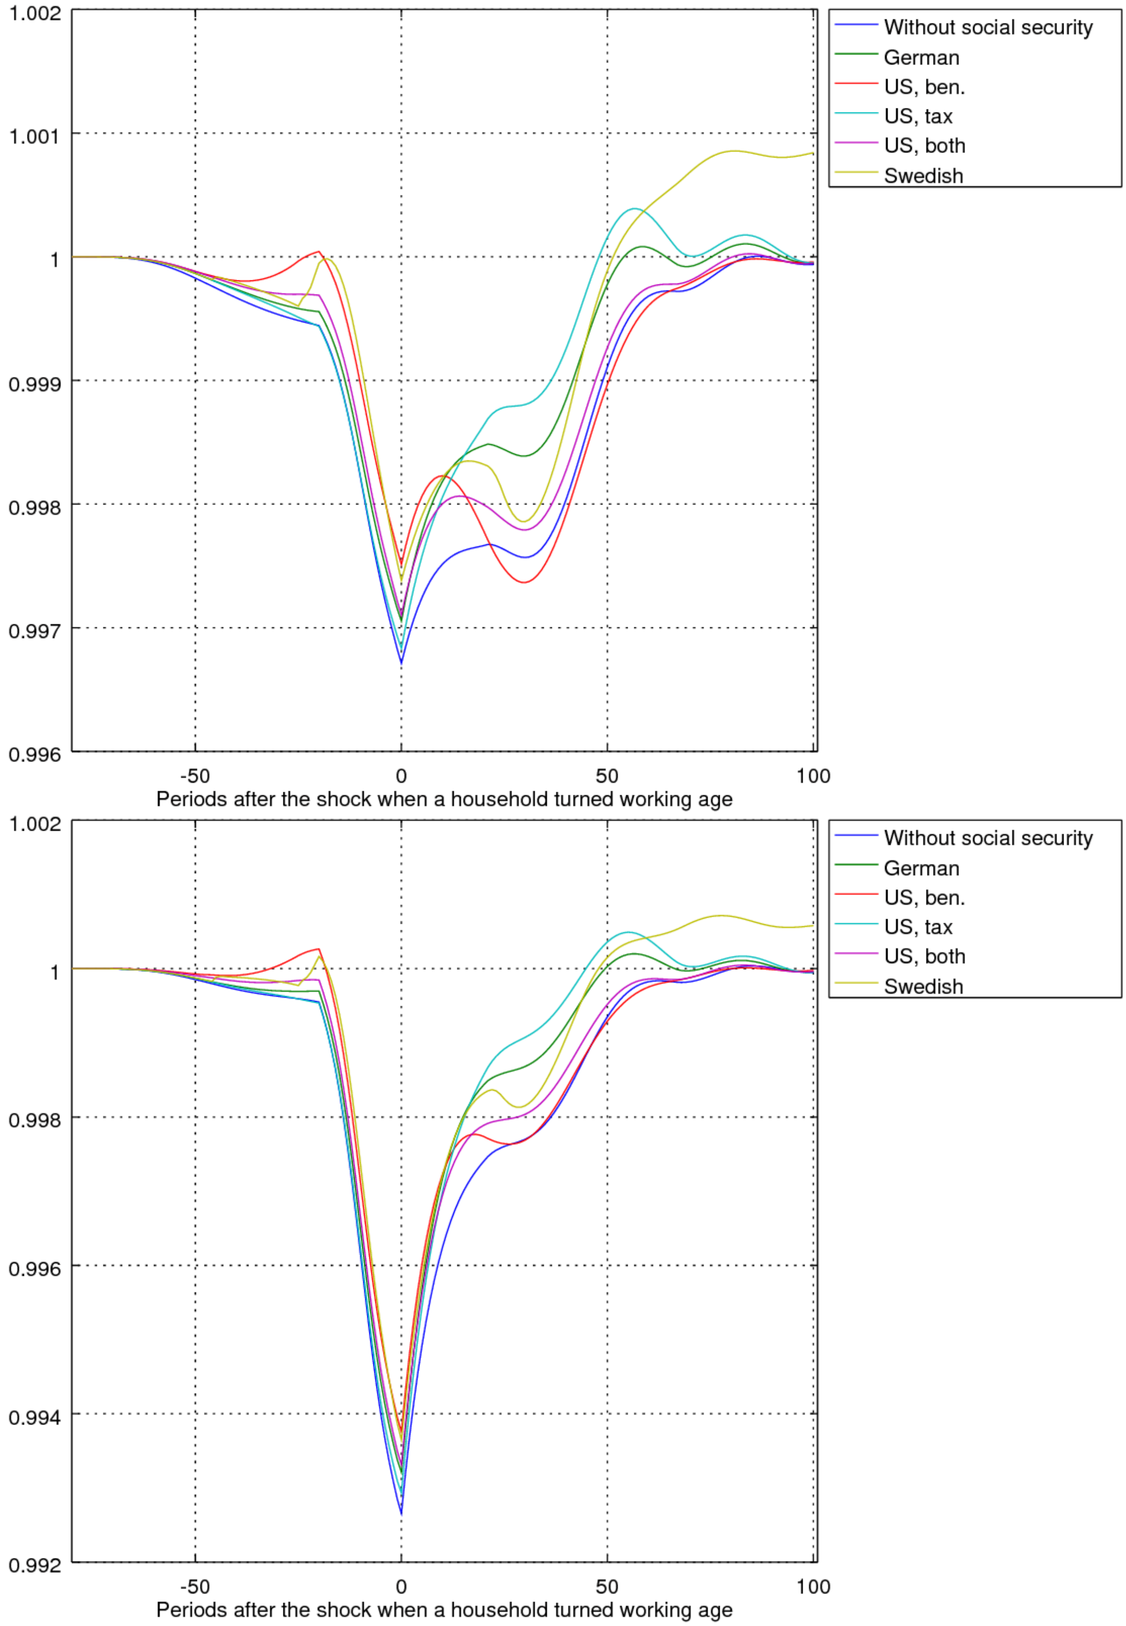

Figure 18. Migration Shock: Wealth Equivalents for γ = 0.35 (upper), 0.7 (lower)

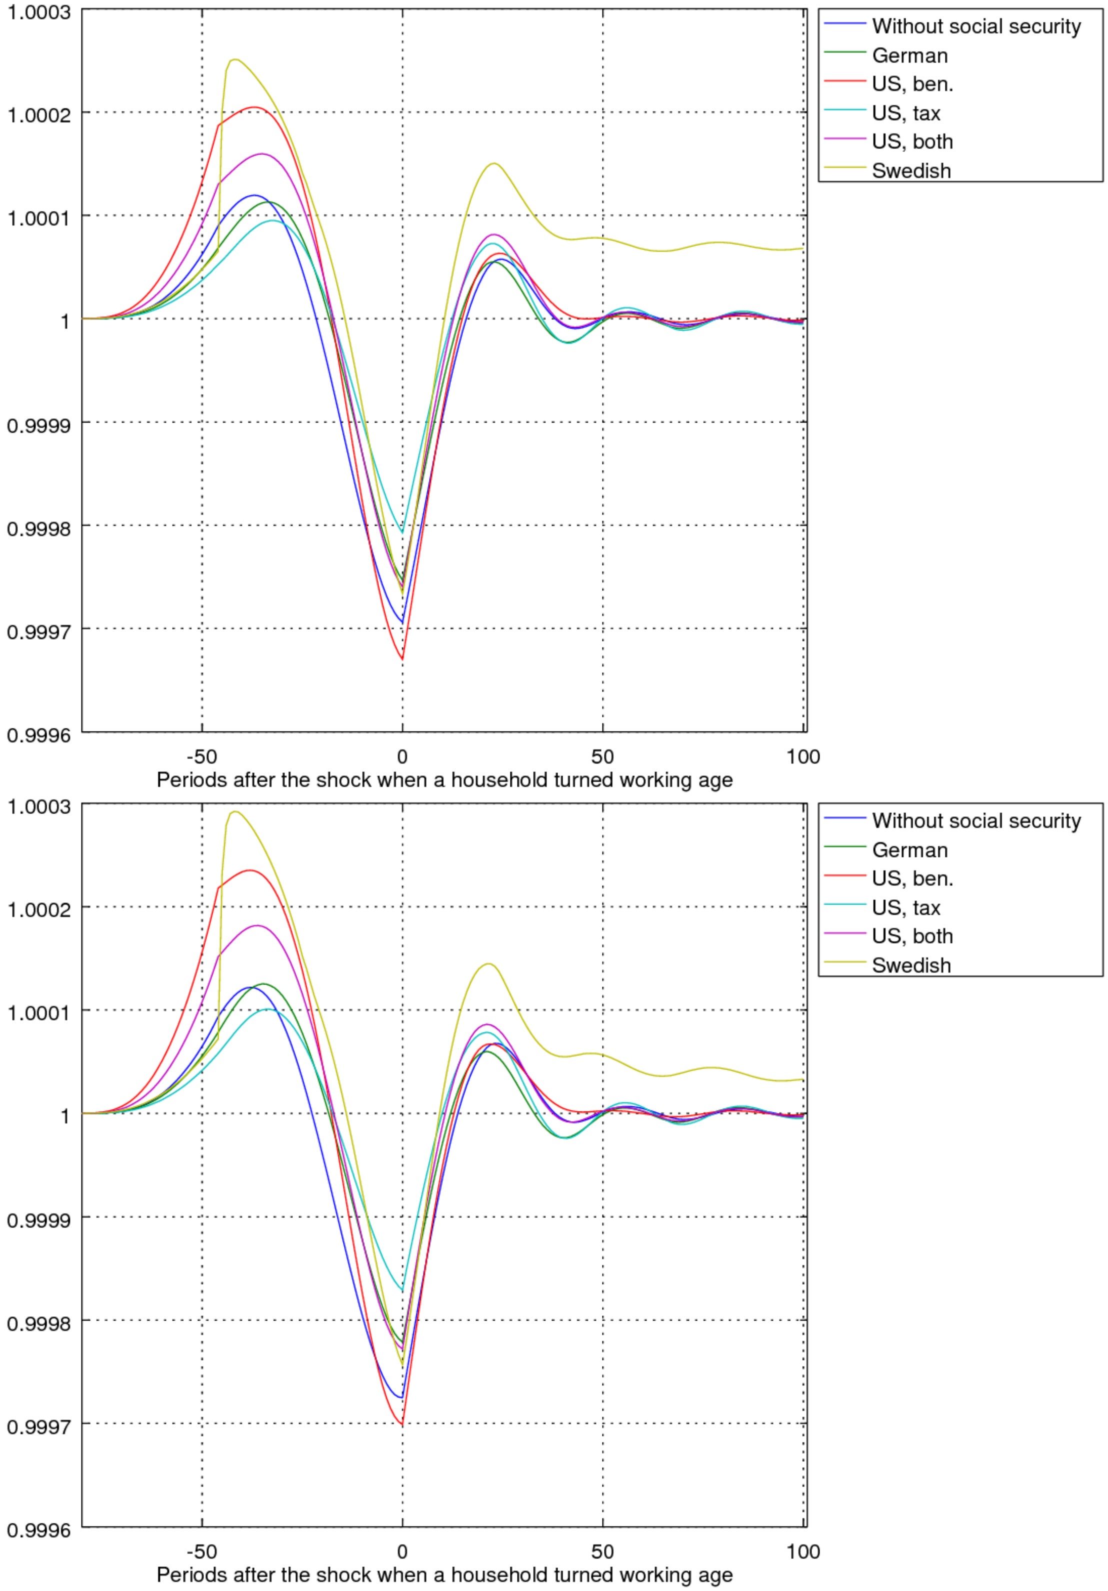

Figure 19. Productivity Shock: Wealth Equivalents for Closed & Open Economies

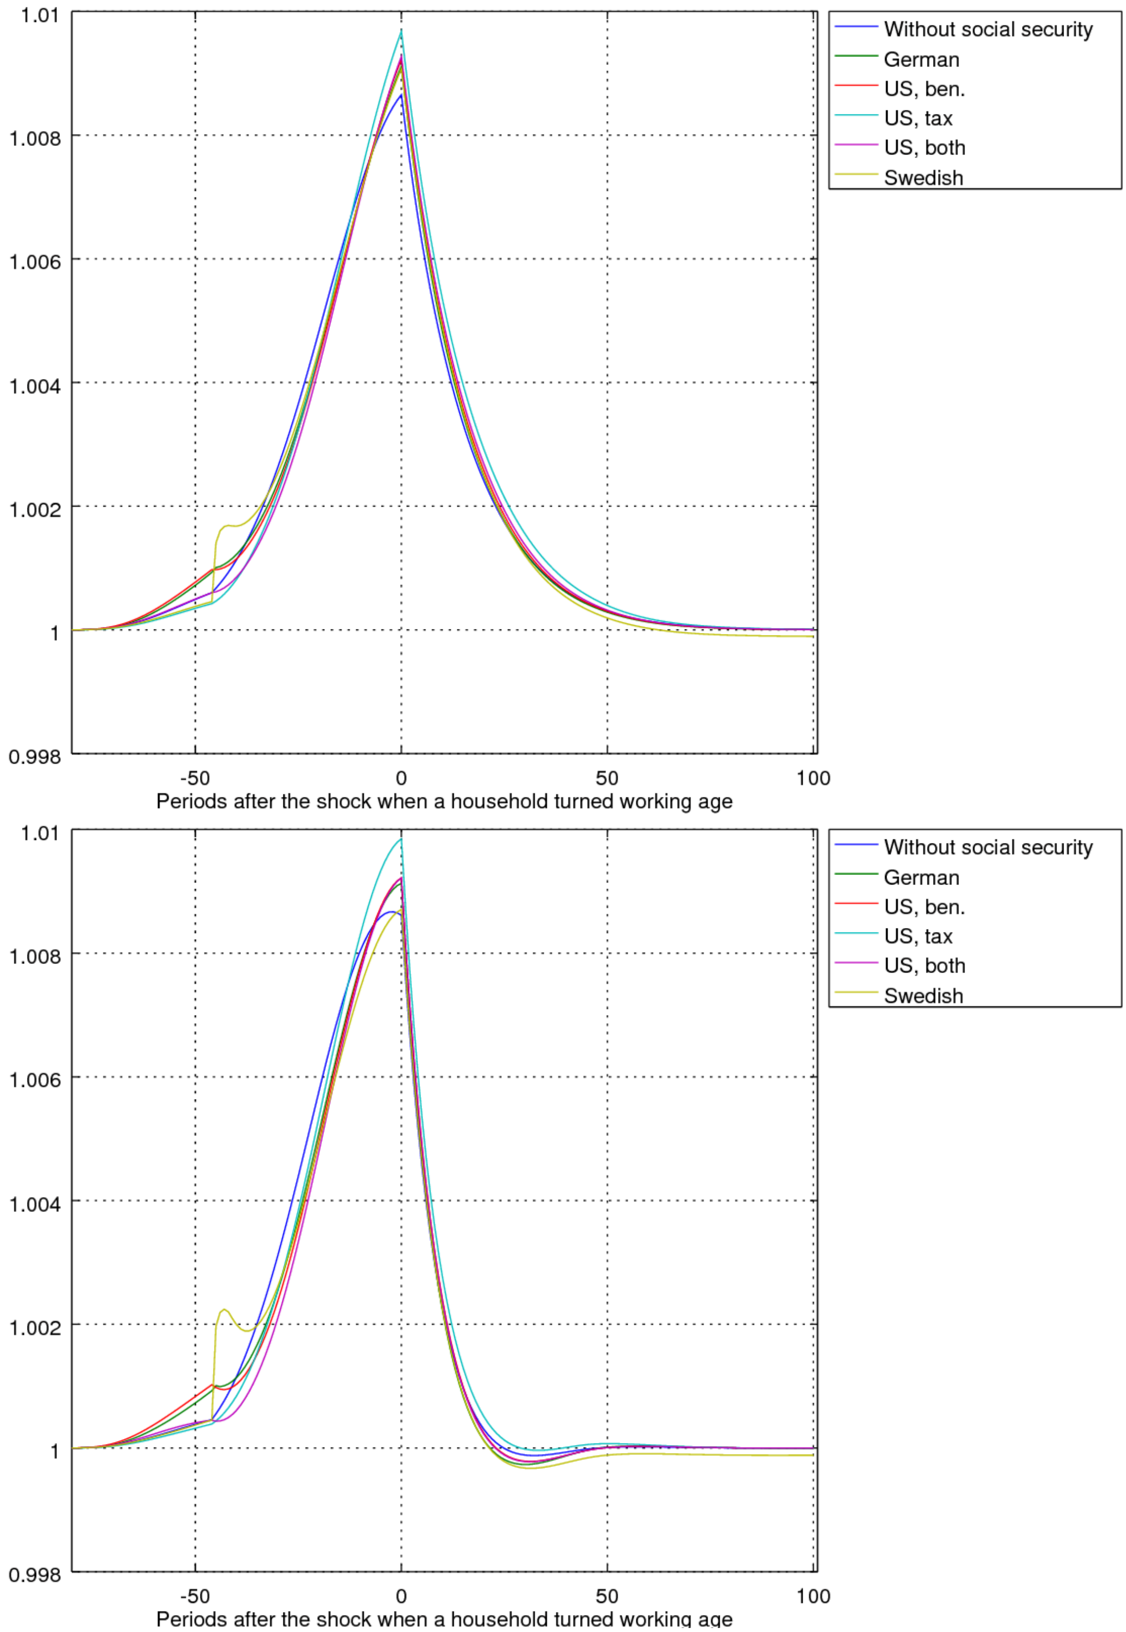

Figure 20. Mortality Shock: Wealth Equivalents for Closed & Open Economies

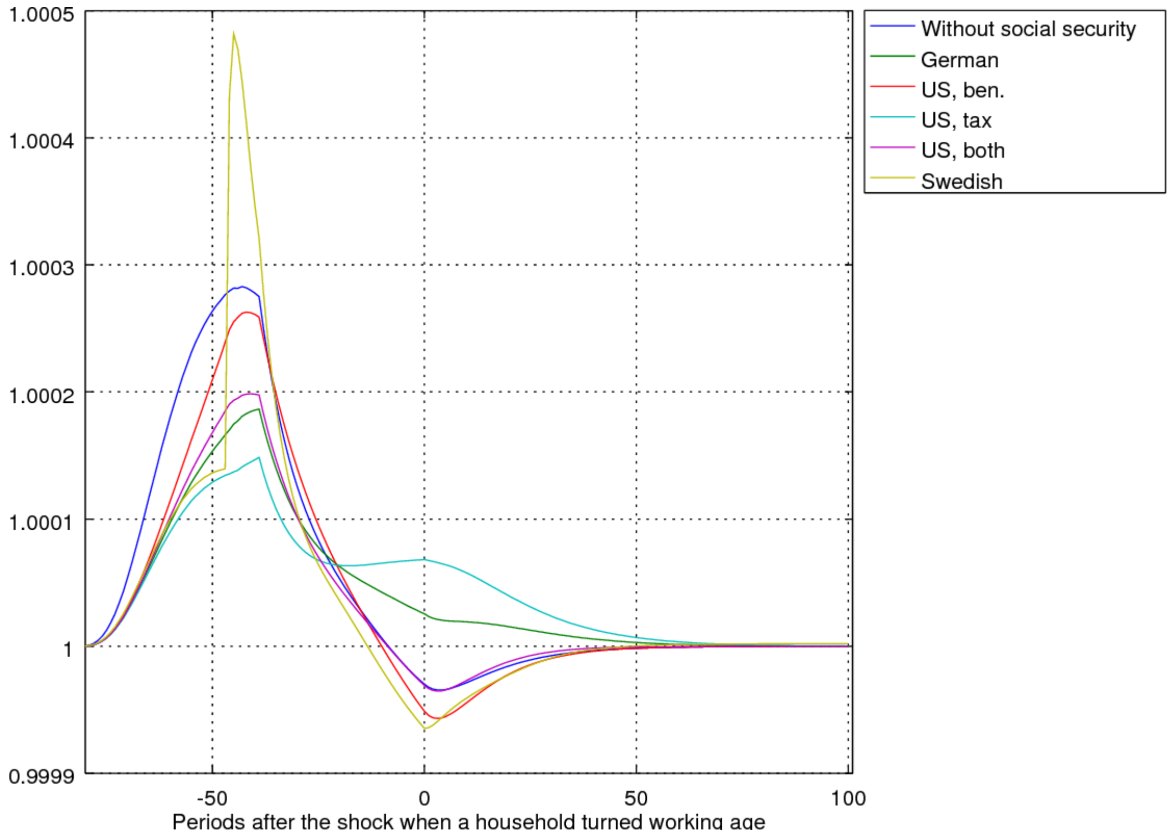

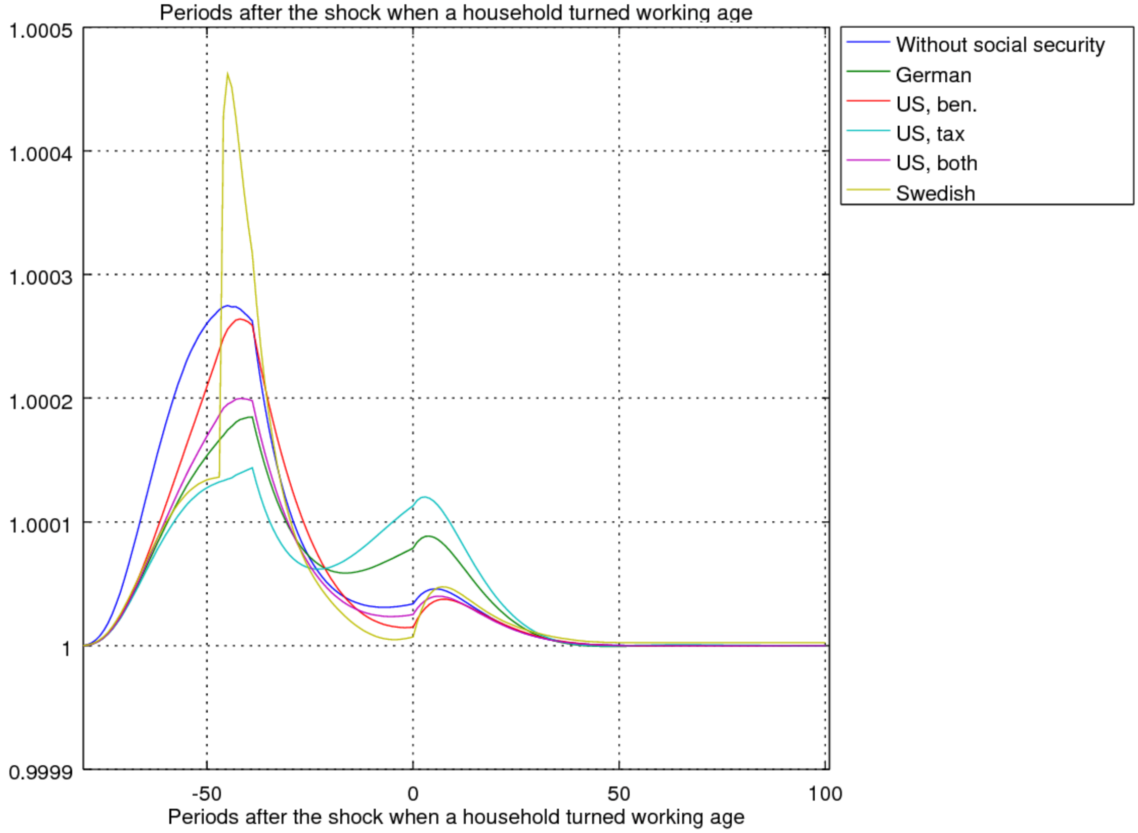

Figure 21. Fertility Shock: Wealth Equivalents for Closed & Open Economies

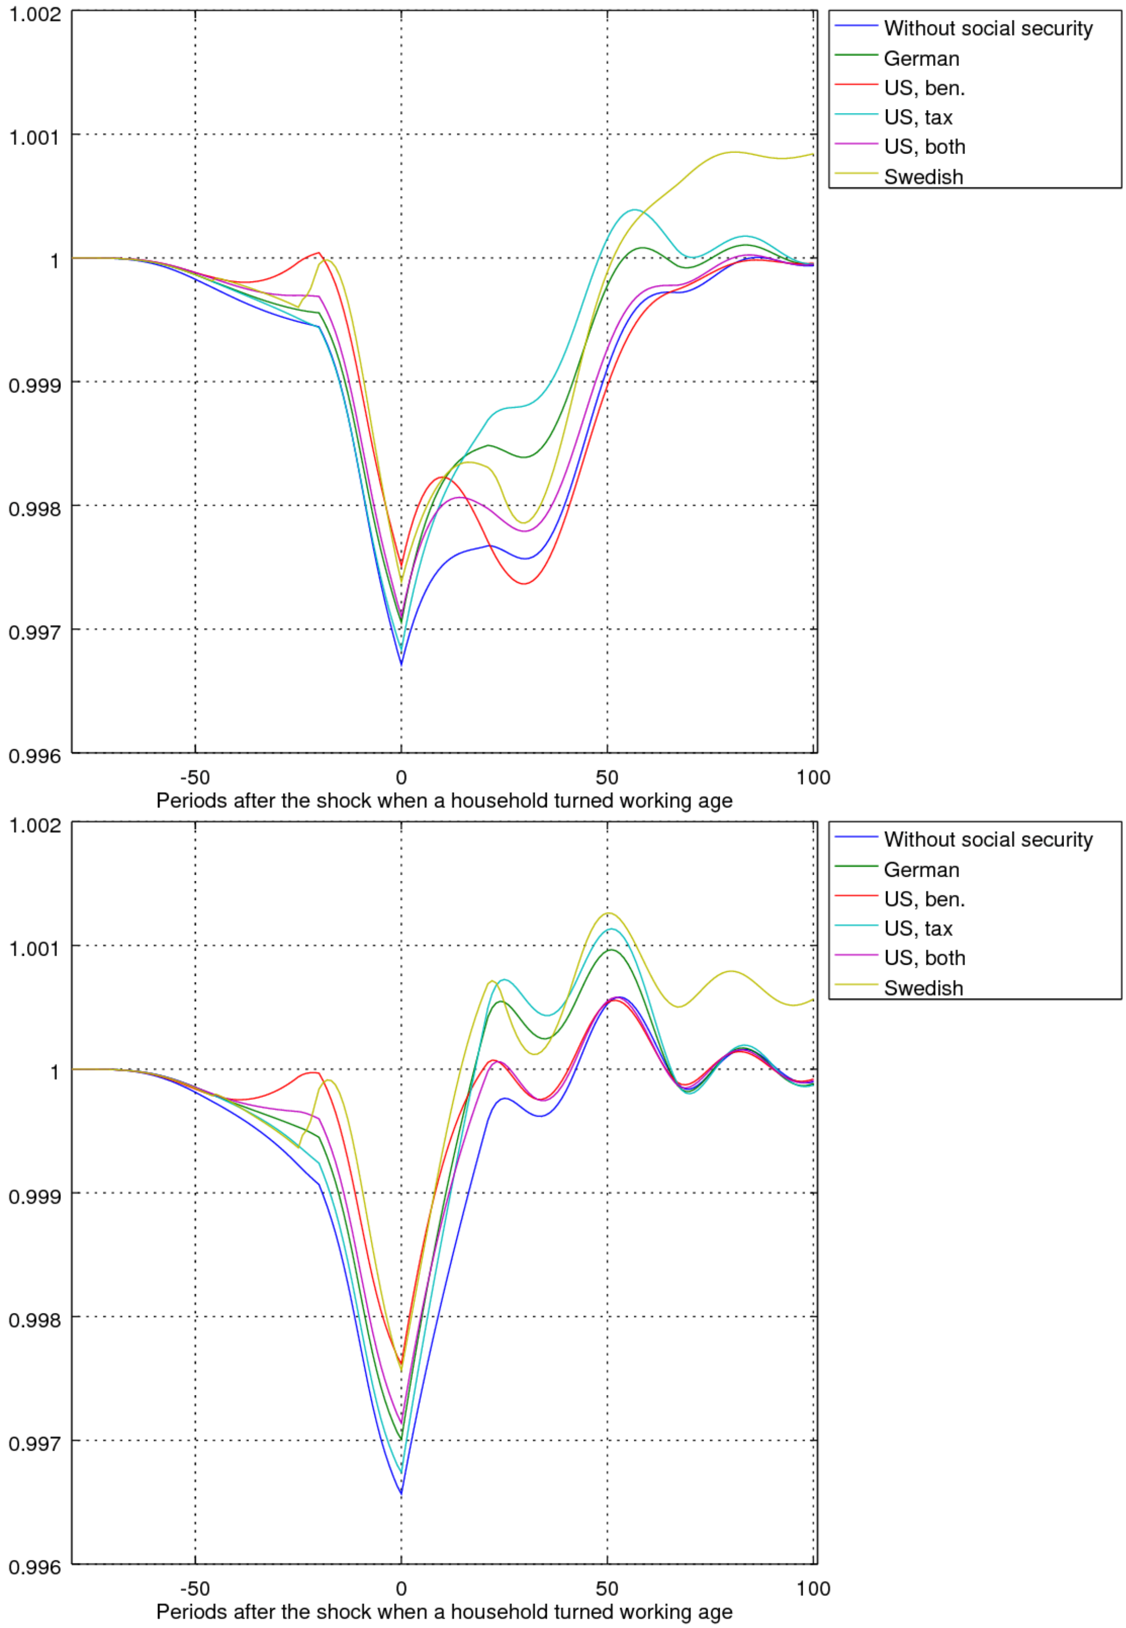

Figure 22. Migration Shock: Wealth Equivalents for Closed & Open Economies)