

Low fertility and long run growth in an economy with a large public sector

Jovan Zamac, Daniel Hallberg and Thomas Lindh

Institute for Futures Studies, Box 591, 101 31 Stockholm

Corr: thomas.lindh@framtidssstudier.se, office +46-8-401216, fax +46-8-245014

Abstract

There is plenty of evidence that growth has a negative relation to fertility and dependency ratios. Recently it has been suggested that low fertility countries may be caught in a trap that is hard to get out of. One important mechanism in such a trap would be social interaction and its effect on the ideal family size. Such social interaction mechanisms are hard to capture in formal models, therefore we use an agent based simulation model to investigate the issue. In our experimental setup a stable growth and population path is provoked into a fertility trap by rising relative child costs linked to positive growth. Even rather large increases in child benefits are then insufficient to get out of the trap. However, the small number of children temporarily enables the economy to grow faster for several decades. Removing the adaptation of social norms turns out to disarm the trap.

Keywords: low fertility trap, social norms relative income, economic growth

Manuscript presented at the European Population Conference 2008 in Barcelona.

Introduction

There is by now plenty of evidence that growth has a negative relation to fertility rates and dependency ratios. Kelley and Schmidt (2005) summarize much of the evidence concerning the importance of demographic factors for growth in both developing and developed countries. Lindh and Malmberg (2007) estimate a demographically based forecasting model allowing for changing economic impact of the age distribution with rising longevity. The negative impact of large cohorts of children is a robust feature. In Barro-type cross-country growth regressions high fertility has also been established as a substantial negative factor. This is not really news but confirmation of for example the shift-share analysis of Anne O Krueger (1968) making the point that a very large part of the difference in production levels between developing and developed countries can be explained by demography and education.

It is intuitively rather obvious that an economy where almost half of the population is below 15, as is the case in many African countries, has to carry a large burden of supporting and educating the young, a burden of which only a tiny piece will appear as value-added in the national accounts. As Bloom, Canning and Sevilla (2003) emphasize the demographic dividend from falling fertility appears as the working population starts to grow faster than the dependent population. Mason and Lee (2007) stresses the possibility of a second demographic dividend as middle aged people start saving for retirement and thereby contributes to increasing capital resources.

Because growth can be accelerated by lowering fertility it becomes theoretically possible to compensate rising elderly dependency rates with decreasing fertility. Indeed several developed countries both in Europe and Asia are experiencing very low fertility. While only China has a deliberate one-child policy, fertility inadvertently seems to be lower than in China in both South Korea and Japan. Since fertility lower than replacement rates will speed up aging in the population this is ultimately not a sustainable equilibrium. Sooner or later the support for the elderly in the economically active population will become insufficient. That may well take a long time though and in the meantime resources can be shifted from reproduction to support for the elderly.

Recently it has been put forward, e.g. Lutz et al. (2006), that countries with very low fertility (below say 1.5) may be caught in a low fertility trap (the Low Fertility Trap Hypothesis, LFTH). This builds on the observation by inter alia McDonald (2005) that a recovery towards replacement levels seems to become increasingly difficult to achieve in countries with very low fertility. The mechanism suggested by Lutz et al. builds on three different feedbacks. One is the demographic inertia of a baby bust where the small cohorts will have fewer babies simply because they are fewer in fertile ages. The second feed-back is hypothesized to work through social norms making fewer children more desirable as the prevalence of children in society becomes rarer. This social interaction and its effect on the ideal family size work as a negative feedback reinforcing low fertility standards. Thirdly Lutz et al. adds the socio-economic relative income effect of Easterlin (1961), where the aspired level of living is determined by consumption

standards in the parental home acting as a constraint on the number of desired children when it becomes harder to achieve these aspirations (also cf Macunovich 1998). Such social interaction mechanisms are hard to capture in a standard economic model framework. For this reason we propose to use an agent based simulation model to investigate the economic consequences of low fertility and its feedback on the fertility decision.

Our purpose is to use a simple agent-based framework model as an experimental tool to investigate under what circumstances a low-fertility trap may be likely to appear. The model is built around the tension between a long education interfering with the prime fertility period of females and the need for such education in order to satisfy consumption aspirations arising from the relative income effect.

The crucial element in an agent based simulation model is the rules of thumb that agents use for making decisions. These rules can incorporate both an economic dimension and a social interaction dimension which are important for the low fertility trap. Individual micro behavior results in a macro outcome which in turns feeds back into individual decisions, and hence a micro-macro interdependence is obtained which cannot be modeled in more traditional microsimulation models.

By calibration to Swedish micro and macro data¹ the simulation model offer an experimental laboratory to test different theoretical mechanisms and their implications for the balance between current benefits from low cohort fertility and the losses in terms of future shrinkage of the tax base and the growth potential. Thus our base scenario approximately reproduces the natural reproduction features of the Swedish population during the 20th century.

We then experiment with alternative scenarios focused on the relative income mechanism in order to spring a low-fertility trap. It turns out that introducing a mechanism that increases the relative cost of children versus consumption has the potential to set off a low-fertility trap. The growth rate in GDP per capita increases substantially and for a long time before the elderly dependency burden makes the system fiscally unsustainable. It takes a very large and persistent increase in child benefits to counteract this and climb out of the low fertility trap. The crucial element of the fertility trap turns out to be the social norms system. With fixed social norms the increase in relative costs of children does not provoke a fertility trap.

The disposition of the paper has to be parsimonious with details about the model, which are briefly described in the appendix. In the next section we first discuss the background and theoretical starting point more completely. Thereafter a short account of the model structure and the crucial mechanisms in this context follows. In the third section, we analyze the results of alternative scenarios and discuss the potential implications. Finally we conclude our argument by discussing how the results can be used to focus further work around growth and fertility.

¹ We have also used some data from the Swedish National Transfer Accounts. See Lee and Mason (2004) about National Transfer Accounts that tracks intergenerational flows in an economy.

Background and theoretical starting point

Negative momentum from low birth rates that decreases the future fertile population is an obvious demographic feedback from low fertility. Ideal family size as measured in attitude surveys shows a downward trend in many countries. Whereas ideal family size still is higher than actual total fertility rates (TFR) in Europe, there are some countries e.g. Austria and Germany where ideal family size now is well below replacement rates. There seems to be a downward trend of relative income in many industrialized countries (as far as it can be correctly measured) where it is claimed that the next generation will fare worse than their parents. All this seems to coincide with downward trends in cohort fertility and definitely with postponement of child bearing. But will these factors combine into a mutually reinforcing feedback circle as the LFTH suggests? There are three apparent routes of escape to consider. First, extensive immigration may increase the fertile population and reverse the negative momentum. Second, ideological pressure could be mobilized to favor ideals of replacement fertility. Third, social policies can be implemented to relieve the relative income pressure. Our focus will be on the third escape route although we also experiment to see the potential of the second. The first with immigration cannot be implemented in the current version of the model.

The social policy framework differs quite a lot across the countries that now are experiencing negative fertility trends. While the fertility response to social policies seems to achieve little or nothing in raising fertility in most industrialized democracies, studies based on specific countries, although a small group in number, give a somewhat different picture (Caldwell, Caldwell and McDonald 2002). It has been argued that the Scandinavian emphasis on policies favoring dual-earner families explains why the downward fertility trend is much more damped in these countries (Ferrarini 2003). The Swedish case is then very relevant to study as a potential model that refutes the necessity of the LFTH. Sweden has had a very stable cohort fertility rate (CFR) around 2 throughout for all cohorts born in the 20th century where we can observe the completed fertility, i.e. up to the early 1960s. The TFR, however, has fluctuated between 2.1 and 1.5 in the postwar period. Björklund (2006) thinks younger cohorts in Sweden will catch up in their lagging CFR, and indeed birth rates have been picking up and may be headed for a new baby boom. Still it seems to be a fact that the relative income of young adults in Sweden have been falling over the same period. Norms regarding the desired family size have been very stable around 2 as measured in attitude surveys. At the same time the average age of the mother at first birth has risen steadily (24 years 1970, 29 years 2005) substantially increasing the probability of not achieving desired fertility.

Austria and China are two countries that this can be contrasted with. Austria is in many respects similar to Sweden (small industrialized country with an extensive welfare state, fairly similar proportion of immigrants), but with significantly different family policies and lower norms for desired fertility. Both TFR and CFR have shown a strong downward trend and Austria may well be an example where the low-fertility trap has been set off. China on the other hand has had a very fast decrease in TFR enforced by a rather draconian one-child policy, but still TFR based on the 2000 census is estimated to 1.8 (Guo and Chen 2007) not very different from current Swedish TFR at 1.808. South Korea on the other hand have strongly pro-natalist policies that have failed to turn the trend of a

steadily falling TFR from 1.72 in 2000 to an expected 1.28 in 2007. Also Japan (TFR 1.23) has a similar but more drawn out experience.

Thus, one may hypothesize that family policy is a crucial factor interfering with the mechanisms of the LFTH both by relieving or amplifying economic pressures and maybe also by a strong influence on social norms (or family policy being designed to conform to social norms). However, as the South Korean and Japanese example shows, the design of pro-natalist social policies may be hard to implement efficiently in order to actually achieve any results.

As has been argued in the debate around inequality and its relation to economic growth (Aghion et al. 1999) the accumulation of human capital may well suffer from credit market imperfections since future human capital cannot be used as collateral for loans. Thus public intervention becomes necessary in order to resolve this. And indeed, it is a fact that a very substantial part of basic education is being provided through public subsidies or direct intervention, even in many cases made compulsory. Public education will act to reduce the growth cost of inequality in terms of reduced human capital accumulation. At the same time it tends in the long run also to reduce inequality as the returns to education accrue to the individual. The virtuous growth circle arising from this mechanism can be disrupted by increasing dependency burdens. In an ageing economy this leads to conflict between the needs of the elderly versus the needs of the young with the potential to bring about a low-fertility trap that accelerates population ageing.

Our investigation touches upon many other research fields. In a wider perspective it directly concerns aging research and the financing of elderly welfare and a literature about pension systems that is too large to survey here. More specifically it connects to the literature on generational accounting (Auerbach et al. 1992) with the issue of generational equity in terms of both income and life prospects. More directly it relates to demographic research concerning the “second demographic transition” (van de Kaa 1987), comparative welfare research on family policy effects on fertility (Ferrarini 2003) and economic research on cohort-size effects (Macunovich 2002) as well as gender equity issues (Gustafsson 2001).

There is so far little consensus on whether the demographic transition will end up maintaining world population at approximately stable levels or result in a future shrinking population. A generally accepted theory to generate predictions on fertility is missing today. The LFTH suggested by Lutz et al. (2006) offers a framework to structure further research around the micro-macro feedbacks. There are, however, difficulties since we cannot in general observe these feed-backs in isolation. We are dealing with very long-range processes taking place in a quickly changing social environment. Although OLG models following Allais (1947), Samuelson (1958) and Diamond (1965) have become a standard tool for economists analyzing intergenerational issues as well as general macroeconomics it is still the case that these models quickly turn intractable when the population structure is non-stationary. This has led to a great number of attempts to use simulation in order to investigate their properties (e.g. Kotlikoff et al. 2001) under more

realistic assumptions. Increased realism, however, comes at the price of evermore intransparent calibration and assumptions that are not readily verifiable.

Agent-based modeling offers an alternative where we can observe the aggregate effects of decentralized decision making without very strong assumptions on individual behavior and still maintain a degree of transparency and opportunity to experiment with different mechanisms. Traditional microsimulation models building on estimated micro relations (Klevmarken 2002) have pretty much come to a point where severe problems have arisen concerning their ability to actually improve our understanding of behavioral mechanisms and their repercussions on the economy at large. Agent based simulation modeling has recently been increasingly explored as a more flexible alternative focusing on actual behavior rather than the optimal behavior of individuals in recognition of the fact that even if agents do behave rationally under their respective information sets, statistical methods will not allow us to evaluate the full heterogeneity of individual behavior (Richiardi et al. 2006).

While the flexibility of ABM allows using it as a laboratory to experiment over a wide range of issues its drawback is that it easily tempts the researcher to try to do too much, to keep too many options open, to start playing a Sims game instead of investigating real issues. It is therefore of paramount importance to define the focus of each study rather narrowly. The model that we use in this paper has been developed at the Institute for Future Studies using JAVA programming. The variant used in this paper is adapted to the specific issue of fertility and the mechanisms of the LFTH.

In brief it exploits the interaction of education choice, matching to marry and the fertility decision in order to generate a self-propagating population embedded in an endogenous growth model inspired by the altruistic model of Becker et al (1990). This base version of the model is adapted here to the low-fertility issue by introducing a relative income mechanism.

What Is Agent Based Modeling?

ABM starts from the premise that the "real world" is hardly the work of a central planner, making it conform to rational rules. Rather, the real world is characterized by decentralized, simultaneous interactions between a very large number of different agents, whose decision-making is based on limited rationality, imperfect information, habits and where the local relational context also contributes to those agents' strategies and behaviors.

It has become rather common among economists to want to model micro-macro linkages between individual and aggregate level variables. Most recent attempts have consisted of combining economy-wide Computable General Equilibrium Models (CGE), with microsimulation models (see Davies 2004, for a review). They rely on the classical assumptions (e.g. rationality, perfect foresight, competitive markets, perfect information, market clearing etc.), in order to find an optimal solution or "equilibrium" for aggregate level variables such as total output. CGE cannot account for heterogeneity across households, preferences or technologies; only a few types of representative agents are

assigned the same production or utility functions. This is clearly a simplification which overlooks important variations at the micro level, and more generally makes distributional analysis unfeasible (i.e. how total output and consumption are actually distributed between different agents and what drives these differences).

Microsimulation models on the other hand, are mostly used to study distributional effects e.g. of tax and benefit systems, at the micro level, including (in the case of dynamic microsimulation models) projections over the individual agent's entire life cycle (including behavioral responses e.g. labor supply, fertility choice, education etc). Usually built on household survey data (or other micro-level data), they allow access to detailed information e.g. about individuals' income sources, areas of residence, past employment history etc., but they cannot deal with modeling the monetary side of the economy or with the inclusion of structural macro features and aggregate feed-backs, which therefore have to be assumed as exogenous.

In practice, integrated macro-micro models suffer from difficult implementation, mostly due to a trade off between adding model complexity and finding solutions which can be handled by standard computational tools. ABMs represent a further step in the development of dynamic macro-microsimulation modeling, as they avail themselves of modern computing developments (e.g. object-oriented programming languages) to simulate complex interactions simultaneously, and how these interactions evolve in time through the accumulation of new information, with no need to have two separate converging models (e.g. one micro and one macro), nor to have convergence to an equilibrium solution at all.

The principle behind ABMs is that of multiple interacting agents who are *goal directed* (e.g. preserving a certain aspired consumption level in our case), and who *try to control their environment*, in a decentralized (i.e. non-coordinated, non-centrally planned) system. ABMs do not assume rationality nor the existence of a pre-defined equilibrium outcome. Agents might behave in sub-optimal ways, but they can gradually learn from their experiences and adjust their behavior to the neighboring environment.

The first attempt to apply ABM to the social sciences is considered to be T. Schelling's "Models of segregation" (1969). Using JAVA or similar object-oriented programming languages, ABM agents are usually implemented in software as objects i.e. computational entities that have initial states (e.g. sex, age), are able to perform some pre-specified action or method, can communicate or share information with others, pass on or even inherit characteristics or behavioral rules.

Running an ABM simply means instantiating an agent population, and let it run forward in time - executing it, rather than solving it. The outcomes of agents' interactions can be observed at any given time by the modeler who only needs to specify some initial behavioral algorithms (the equivalent of classical preferences) and initial conditions for his agents and their environment. These agents can represent people (say consumers, sellers, or voters), but they can also represent social groupings such as families, firms, communities, government agencies and nations.

In our application we view the model as an experimental laboratory for testing mechanisms in a more complex setting than analytical modeling allows yet more transparent and subject to experimental control. This makes inference possible regarding causal mechanisms that cannot be gleaned from econometric estimation on real world data. This does not replace other modes of analysis but is a complement for testing the logic of economic mechanisms in a more complex setting that can to some extent be validated against real world data.

In order to interpret and aid this purpose the model has been calibrated to Swedish micro and macro data and reproduce some of the institutional features of the Swedish welfare system, while other technological and economic details have been suppressed.

The IFSIM model

The IFSIM model, as we call it, consists of a small number of interacting modules. Due to space considerations we only give a brief overview of some features that are important for this paper.

Demographic module

Individuals born outside the model come directly from the initial data set.² Artificial individuals are born inside the model from the model's fertility module at age zero. The maximum age an individual can reach is 110.

New individuals will have a mother and a father and will inherit some of its parents' characteristics. They will from birth belong to a network group; which network group they belong to is at first determined by their parents' characteristics. The networks are segmented first by age groups and then within each age group there is a spatial dimension. A network group contains all those individuals to whom the individual is close. We follow Billari et al. (2006) in defining social "closeness" as a spatial area representing the individual's scope of interaction, by age group. Specifically, agents are arranged along the surface of an imaginary cylinder, whose vertical length is broken into as many segments as there are age groups in the model (at present there are 8 age groups, from age 0 to 110). Individuals will migrate between network groups as they age, and two individuals that belong to the same network group at one age, may belong to different groups at later ages. Over time as an individual age she may leave the parents' network group.

A woman can reproduce between ages 20 and 40. The timing of giving birth is determined by a decision function taking account of a) the consumption cost of a new family member, b) the desired number of children (dependent on number of siblings of the parents), c) the social influence of how many children there are in their current

² The data comes from the Swedish micro data set HUS (See Klevmarken and Olovsson 1993 and Flood et al. 1996), which is a representative sample with about 3000 individuals which we scale to obtain out wanted number of individuals in the initial year of simulation. In the first period of the model we also introduce some initial aggregate variables (like, e.g., mortality rates) from Statistics Sweden.

network, and d) the *aspired consumption level* that has been inherited from the parental home in relation to the *projected consumption* for a new child.

As postulated by Easterlin this aspired consumption level is a norm that is formed during youth and which has the consumption level of ones home of origin as a reference point. “Youth” is in the model set to the age of 10. The aspired income to achieve this goal is relative to the income distribution and corresponds to a certain target percentile in the income distribution weighted with an absolute standard-of-living target. The projected consumption level based on current conditions gives the budget restriction. Therefore the comparison is between consumption when you yourself was 10 years old and when your future child is 10.³ In brief, the desired number of children that the couple strives to obtain may not be reached due to their economic outcome and the social influence from their network group.

Once a child is born, its mother is on parental leave for three years before returning to her previous labor market status.

Starting from age 18, individuals living with their parents may start to leave the parental household and set up a household of their own. The decision to leave home is modeled as an exponential probability function depending on age and on the share of youngsters living with their mother within the network.

The matching process is assortative such that pairs with similar human capital and belonging to the same network are more likely to create a new married household.

Educational module

When reaching the age of 7, all individuals are universally put into basic schooling up until the age of 16, corresponding to 9 years of compulsory education. If entering secondary school, the individual will stay in school another 3 years. The choice to apply to university is determined by the individual’s prospective earnings compared to her aspired consumption level, and her preferred number of children, which was mentioned above. If secondary education is enough to reach the same per capita equivalent income as their parents’ given that one’s preferred number of children is reached, they will not apply for university. Hence, the educational choice does not depend *directly* on fertility choices. However, it does so indirectly since if the individual projects that she can reach her aspired consumption level (see above) given her preferred number of children without having to invest in education, then she will not choose university.

The university applicants are ranked according to accumulated human capital (see below) such that the ones with the highest human capital are actually accepted by the University. The number of available positions at the university is set to a fixed proportion of the current number of individuals aged 19 to 30. If attending university, the student will be

³ The consumption projection in ten years is estimated from the sample of individuals in the simulation whom at the fertility decision has ten more years of labor market experience. We correct for the fact that these individuals live in households with different age composition of the children.

entitled to a student allowance for the duration of the course (5 years), amounting to a fixed proportion of average earnings.

Modules for human capital formation, the labor market and consumption good production

We postulate a production technology that only depends on human capital (i.e. there is no savings into other types of productive assets). Human capital is the only asset and it is measured by the quality of labor embodied in each worker, for a given amount of schooling, training, work experience etc. The individual human capital level is accumulated and measured as the output from a human capital production function which aims to capture the quality of educational institutions and the ability of the system to transfer knowledge between individuals.

Human capital entails an externality embodied in the private accumulation of human capital. When someone increases their human capital, e.g. through an additional year of schooling or on-the-job training, their action has two distinct effects: (i) increasing their own marginal productivity or wages; (ii) contributing to the total human capital stock, which also enters the human capital production functions of current students. This spill-over effect is the key externality that generates endogenous growth in the model.

There are four main inputs into the production of human capital: (i) innate individual ability; (ii) ability acquired from parental influence and parental own human capital levels; (iii) ability acquired through formal education; (iv) skills and expertise acquired through training on the job. The model updates in every period the human capital value for each individual.

The natal human capital is derived from the natal human capitals of both parents (cf. genetic and inheritable ability), and a random term. In the pre-school periods the human capital production function is a Cobb-Douglas type with the sum of parents' human capital and the child's own capital as input factors.

During schooling the *aggregate level* of human capital among the teachers, scaled by the number of students, enters the human capital production function along with the parents' human capital. Hence teaching capacity depends both on the number of teachers per student and their competence. Teacher density is kept constant so that teaching quality only depends on the level of human capital. The human capital increases for every year in work but proportionately less the older the individual gets, up until the age of 55 when it starts decreasing. The rate of increase is higher the higher the level of education, thus generating steeper wage profiles over life for the higher educated.

Individuals that are not in school, above 16 years of age, enter the labor market. The labor market module determine the individual's labor market status. If someone is in work, we assume that they are in full time work. Those who are designated to become teachers do not participate in the production of consumption goods. In the current version, retirement is deterministic. At age 65, individuals are automatically removed from the labor force. From this point onward, individuals are entitled to a retirement pension.

There is a single sector in this economy producing a single consumption good. The amount of consumption good produced depends entirely on the total human capital level of the labor force (excluding teachers). The production function follows a Cobb-Douglas shape with a time-dependent productivity factor. This is meant to capture the idea that returns to human capital investment are increasing, in line with endogenous growth theory, but after a certain level the situation might reverse since the externality from aggregate human capital follows a logistic curve. There are two input factors in the model; the total human capital of the primary or secondary level of education individuals, and the total human capital of the university (tertiary) individuals.

There are no monetary values in the model so earnings are represented by the share of total output produced going to each worker. The allocation of the produced good to workers is separated into two steps. First, the total produced goods are allocated to the two production factors (non-university and university degree individuals) such that the shares reflect each group's marginal product. Then, within each group, the consumption good units are allocated proportionally to the human capital of the individual. This implies for example that, as the supply of university degree individuals is reduced, their marginal product will increase and thereby increase their share of the produced goods. This will be observed by young individuals who will be more prone to choose university and thus increase the future supply of university degree individuals.

State, tax and benefit systems

Beside individual agents, the model includes an institutional agent which represents "the State" who collects and redistribute resources. First, the State calculates the total expenditure bill, by aggregating the costs of the education, teacher's salaries and student allowance, parental leave subsidies, child allowance and pensions.⁴ Once total expenditures are calculated, the State will adapt the tax system so as to raise sufficient revenues to balance the budget (no debt is allowed). The tax system comprises a state and a local tax. The state tax is a progressive tax paid by the top 20 percent of the income distribution. The local tax paid by everybody with positive earnings is a flat tax rate on earnings. It is residually derived to cover the part of total expenditure not covered by the state tax. The individual income tax will therefore be a combination of both the state tax (if eligible) and the local tax. The individual disposable income is therefore the sum of any earnings, pensions, student or parental allowances, minus the income tax.

Model scenarios

Given the basic model setup above both the demographic and the social norm mechanism are included. The relative income mechanism has been implemented by defining an aspired consumption level based on the household income in the parental home at 10 years of age. See the appendix for the details of this. To decide whether to have children a match must first have been achieved with a partner in the local network. In the next step a decision is taken whether to have a child. First it is determined whether a child can be

⁴ The pension system is modelled according to the new Swedish system except for that we have a fixed retirement age at 65 and that we do not have any funded part. We only model the pay-as-you-go component which in reality comprises about 87 percent of the total public coverage.

afforded at the current income level. If so then if future income expectations allow a new child the aspired consumption level at age 10 procreation is initiated. This is the base scenario which we then subject to changes in order to provoke a fertility trap.

Base scenario

In the base scenario the simulation runs for 317 years. The influence from initial conditions takes about 100 years to vanish. An initialization period of 150 years is therefore disregarded in the analysis. After this period the model stabilizes and roughly reproduces 20th century demographic behavior in Sweden.

Worse economic possibilities at fertile ages compared to the aspired income make potential parents postpone children, in hope of better economic circumstances in the future. As the returns to education increase with education level, high educated parents increase the chance to get their preferred number of children. On the other hand education takes valuable time from their fertile years. Since fecundity decreases with age this increases the risk of not reaching the desired number of children.

Below we briefly present some main features of the base scenario. In Figure 1 we depict the population development in the last 167 years of the simulation (for convenience we label model years starting from the year 0 but note that this is arbitrary). There are short term periods of positive and negative growth, but the long term trend is positive population growth. Over the study period there is an increase from about 11,000 to 15,000 inhabitants.

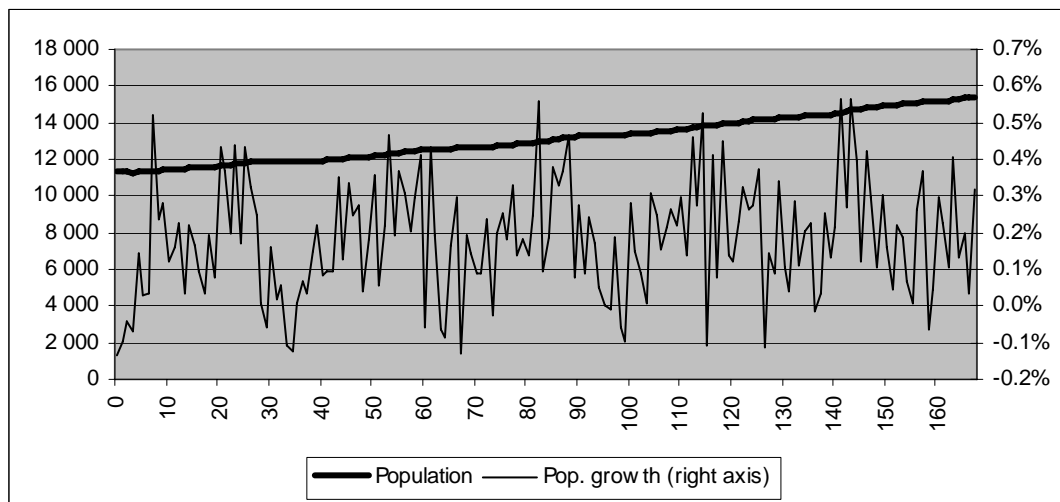


Figure 1 Population development in base scenario

The age-composition in the population is very stable over time, see Figure 2 below, which shows the shares of young (0-20), prime-aged (20-64), retired (65+) and oldest old (80+) by model year. There is an oscillatory pattern (suggesting a saw-toothed age distribution) that reflects influential baby boom cohorts which, to some extent, actually resembles quite well the Swedish demography over the 20th century.

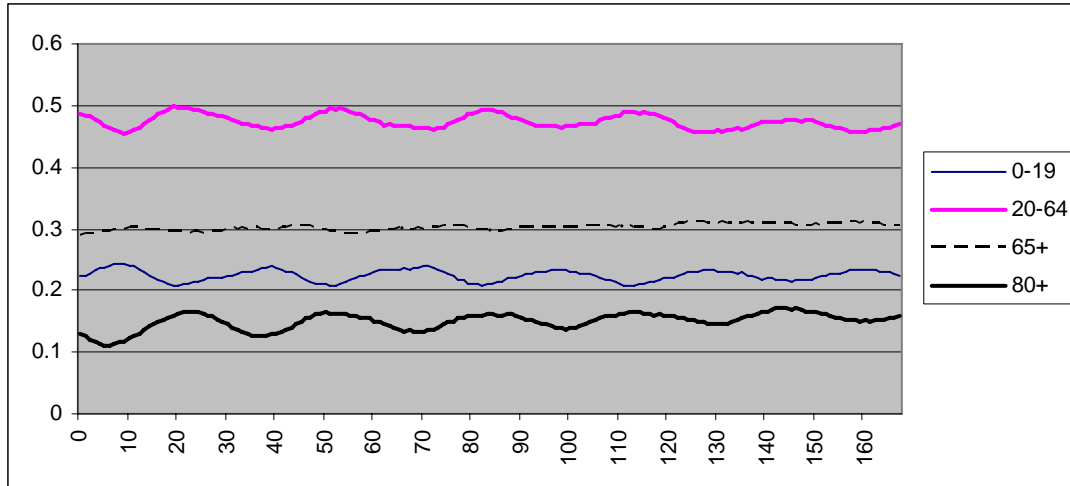


Figure 2 Age structure in the base scenario

There is variation in the CFR over time that explains this pattern. For the whole period CFR is on average above reproduction level 2.15 children per woman. We observe a rather steady pattern but there are some temporary swings in CFR, see Figure 3. We do not have any mechanism that leads to intensified efforts to procreate as fecundity goes down. Intuitively there is no “biological clock” that makes individuals try to catch up in their 30s. Therefore the swings in period fertility due to changing economic conditions mostly carry over to cohort fertility.

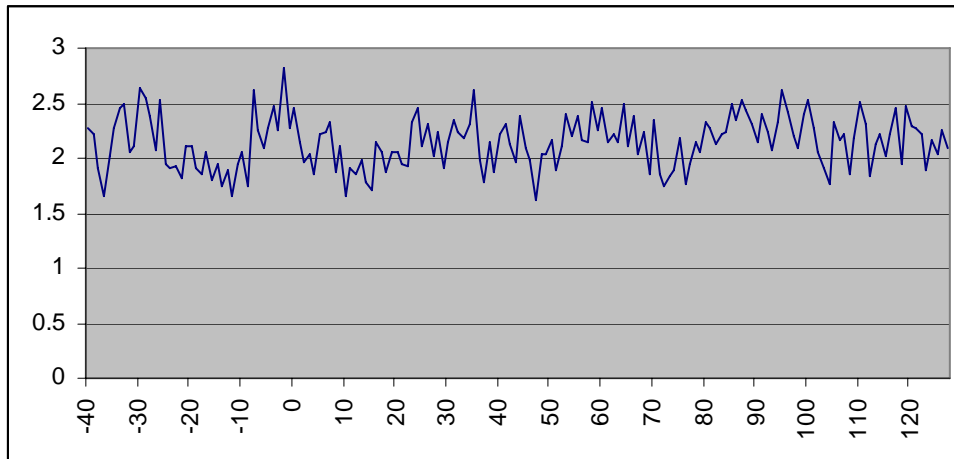


Figure 3 Completed fertility, by birth cohort

As can be seen from Table 1 the age of the mother in general increases with education. Mothers are on average about 27 years old when they get their first child if they have only high school or basic education, while mothers with a university degree are on average about 29 when they have their first child.

Table 1 Average age of mothers when giving birth in the base scenario, by the number of children and education level

<i>Number of children</i>	<i>Basic+High school</i>	<i>University</i>	<i>Total</i>
1	27.4	29.0	28.1
2	30.3	31.2	30.7
3	33.2	33.9	33.5
4	34.9	35.5	35.2
5	36.2	36.6	36.4
6	36.8	37.4	37.1
Total	30.7	31.7	31.2

One can note that university studies are spread over a rather long period in young adulthood as some actually enter at a late age. The typical age to get a university diploma is age 24 (i.e., the youngest possible diploma age), but the rest, about 17 percent, graduate between 25-40. These delays in achieving a diploma will push child birth into later in the fertile period. There is also a small chance that students have a baby during education enrolment. Such events will delay university diploma even more for female students (by three years per baby). A larger effect on high educated may reflect the increased difficulty to combine education and children. At the aggregate level the base scenario data also shows that rising enrolment rates in higher education are linked with lower birth rates. Once enrolment rates drop, individuals achieve their diploma, and birth rates increase.

The development of earnings over time, presented for the years 60-125 in Figure 4, shows strong period effects. These are due to shortage of one education class relative the other creating a wage drift upwards for the scarce labor category. In comparison, most of the period effects are not reflected in the human capital stock, shown in Figure 5 for the same period. The ratio of earnings to human capital hence fluctuates over time. The no education group are children below age 15 who do have human capital but no earnings.

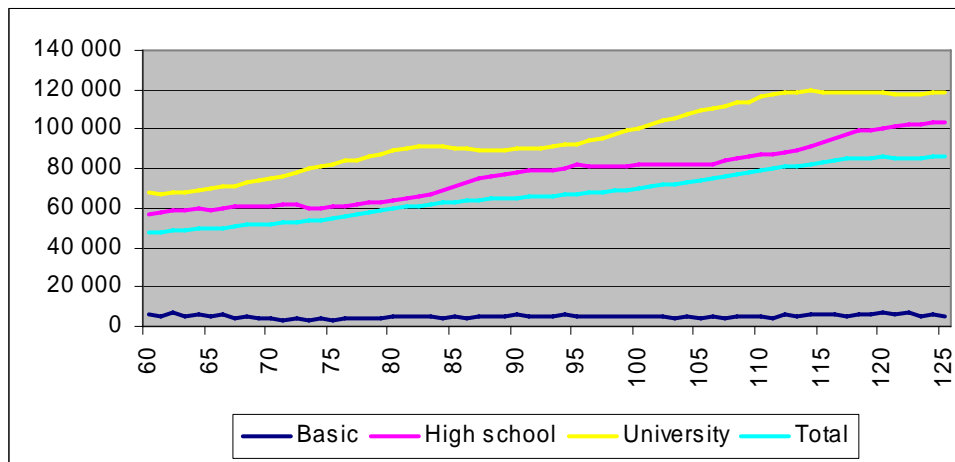


Figure 4 Earnings in model time 60-125, by education, base scenario

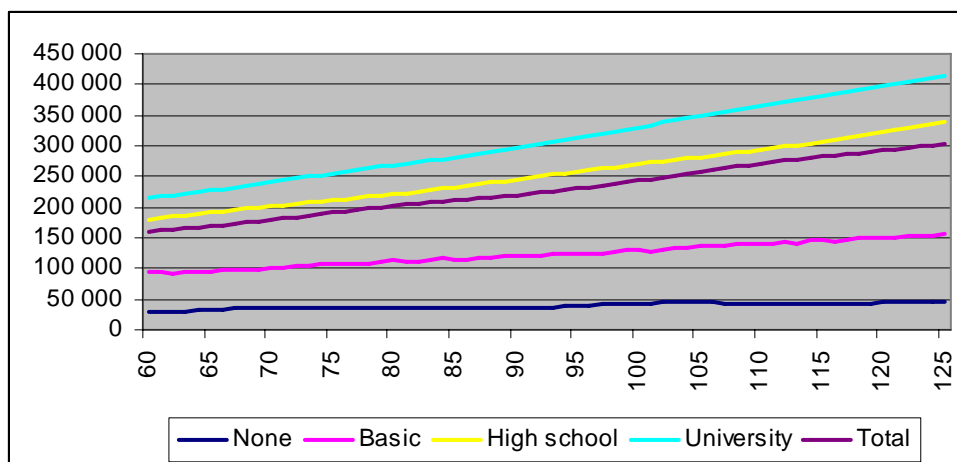


Figure 5 Human capital development, model time 60-125, by education, base scenario

Alternative scenarios with a higher cost of child

In five alternative scenarios we made modifications to the function that determines the cost of children in order to provoke a fertility trap. In the base scenario the expected cost of having a new child is a fixed share of current median income. In the alternative scenarios when the economy grows faster than usual there will be a relative price shift making children costlier in terms of consumption. This will force parents to invest a higher share of their income in the child, if they decide to get one. One rationalization of this upward drift in the relative cost of children is that the opportunity cost for parental time is increasing. Since everybody works full time in our model we cannot implement that directly through labor supply.

Below we describe the five scenarios.

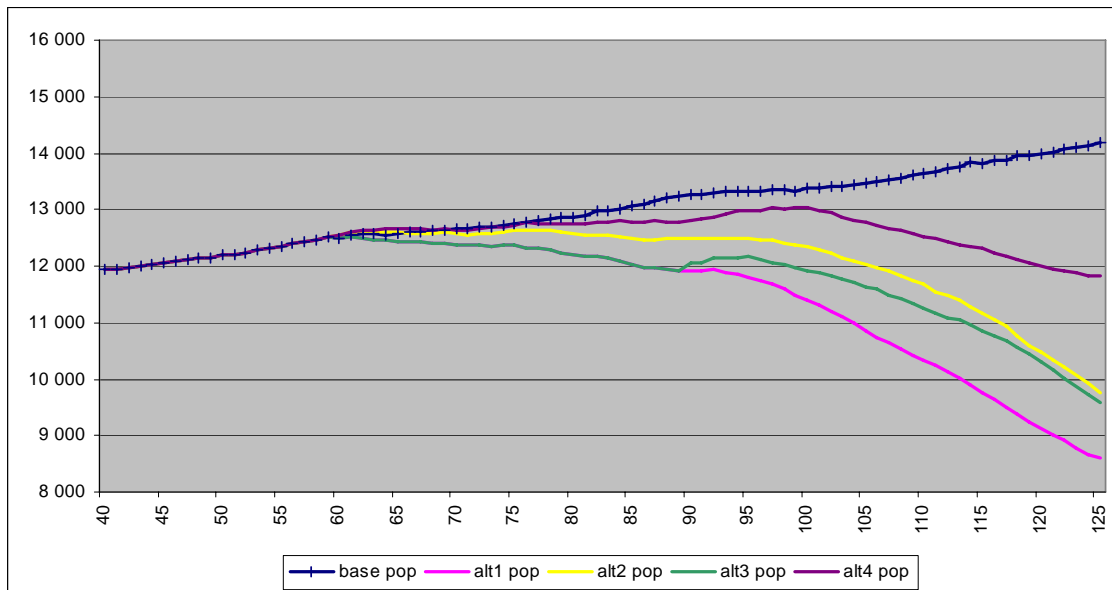
1. In our first alternative scenario, “do nothing”, no policy is introduced to counteract the change in cost of children.
2. In the second scenario, “do something immediately”, a child benefit is introduced in the same year as the relative cost mechanism is introduced. The benefit is given to a child aged 0-17. The benefit per child is set to a fixed 10 percent of average earnings. Explicitly the benefit is given by $b_t = 0.1 * \frac{\sum_i W_{it-1}}{N_{t-1}}$ where W_{it-1} is i :s earnings and N_{t-1} is the number of wage earners, both measured in the previous year $t-1$. The benefit will follow the GDP growth and be independent of the number of children.
3. The third scenario, “do something but maybe too late”, is just like the second one but implements the benefit with a 30 year lag.
4. The fourth alternative scenario, “do something drastic”, implements a strong form of the child benefit immediately like in scenario 2. In this scenario the benefit is

given by $\tilde{b}_t = 0.1 * \frac{\sum_i W_{it-1}}{n_{t-1}}$ where n_{t-1} is the number of children eligible for the benefit. \tilde{b}_t will in general be higher than b_t . In a situation when few children are born, the benefit \tilde{b}_t will rise to counteract the upward drift in the relative cost of children.

5. The final scenario is scenario 1 but turns off the social norm mechanism.

In sum alternative 1 and 3 does nothing or too little, while alternative 2 and 4 implements active policies very quickly. Alternative 5 includes no policy responses but prevent social norms from adapting. This scenario is treated separately below.

The relative cost mechanism is introduced in the model year 60.⁵ One can clearly see that fertility is affected by the introduction of the changed child cost. In the first and third alternative scenarios population actually declines at a fast rate (Figure 6) and is also rapidly ageing (Figure 7) directly after the introduction of the magnified expected child cost. We note that in the first alternative scenario the share 65+ in the population passes 50 percent about 65 years into the future after the shock. As is obvious in Figure 8 the CFR declines and two decades after the introduction of the relative cost mechanism the CFR only rarely pushes over reproduction level. The child benefits only delay the development. A late implementation has only very temporary effects. The relatively large child benefit introduced in scenario 4 is most successful in blocking the dive in fertility. However, 40 years into the future, not even this policy manage to turn the tide as fertility rates falls to low levels.



⁵ All alternative scenarios are programmed in such a way so that perfect replication of the base scenario is attained, up to a point when a change of arbitrary choice (like the change in child cost) is set into play. This means that the initial random component in the model is exactly identical in all scenarios. Thus the alternatives can be interpreted as counter-factuals to the base scenario.

Figure 6 Population in the base and alternative scenarios

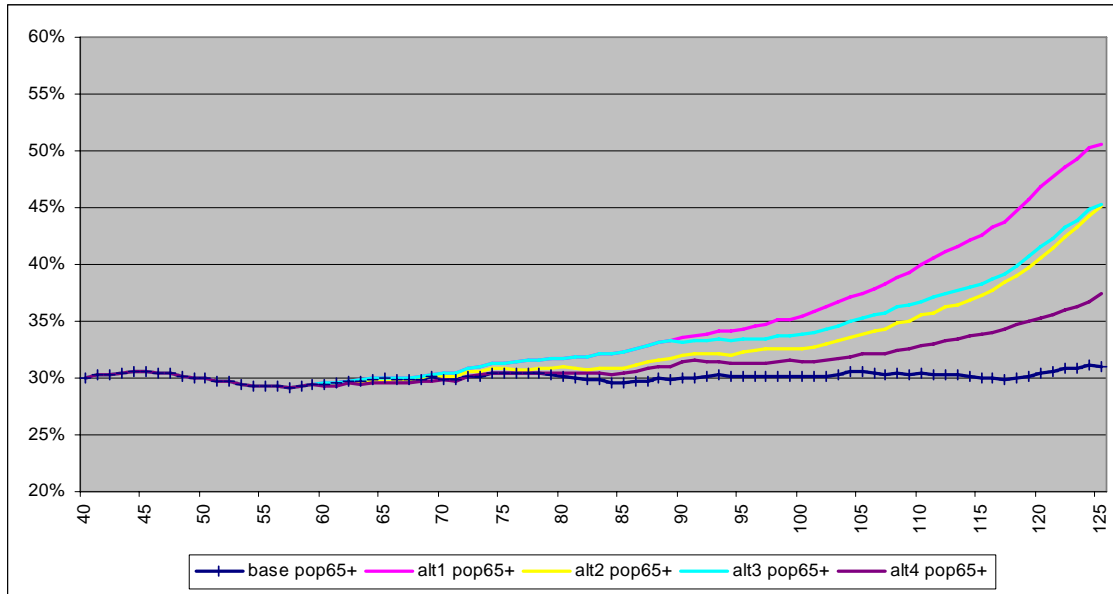


Figure 7 Share of age 65+ in the base and alternative scenarios

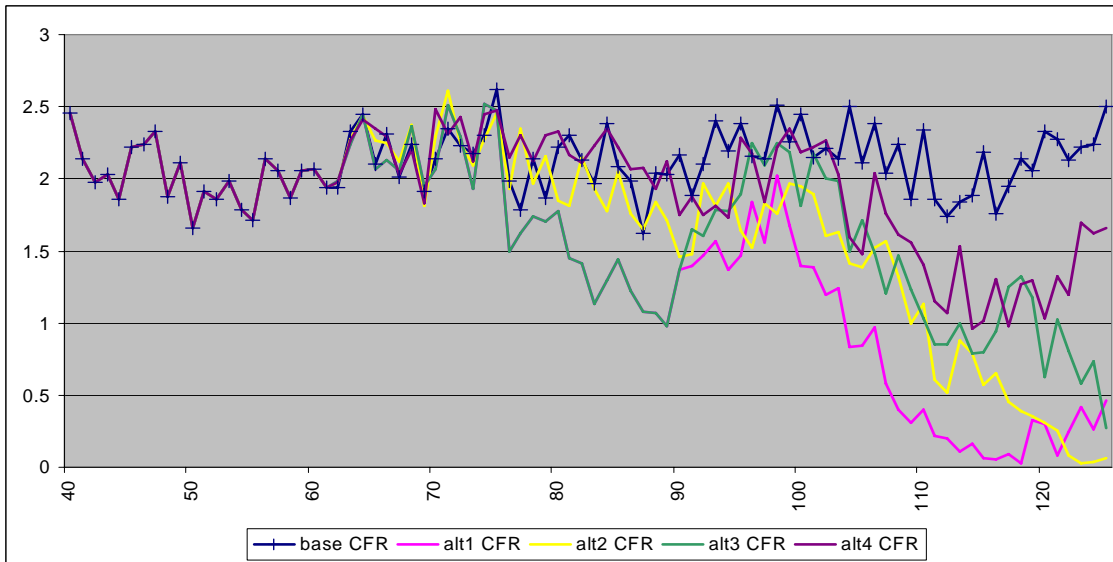


Figure 8 Completed cohort fertility rate (CFR) in base and alternative scenarios, by model time

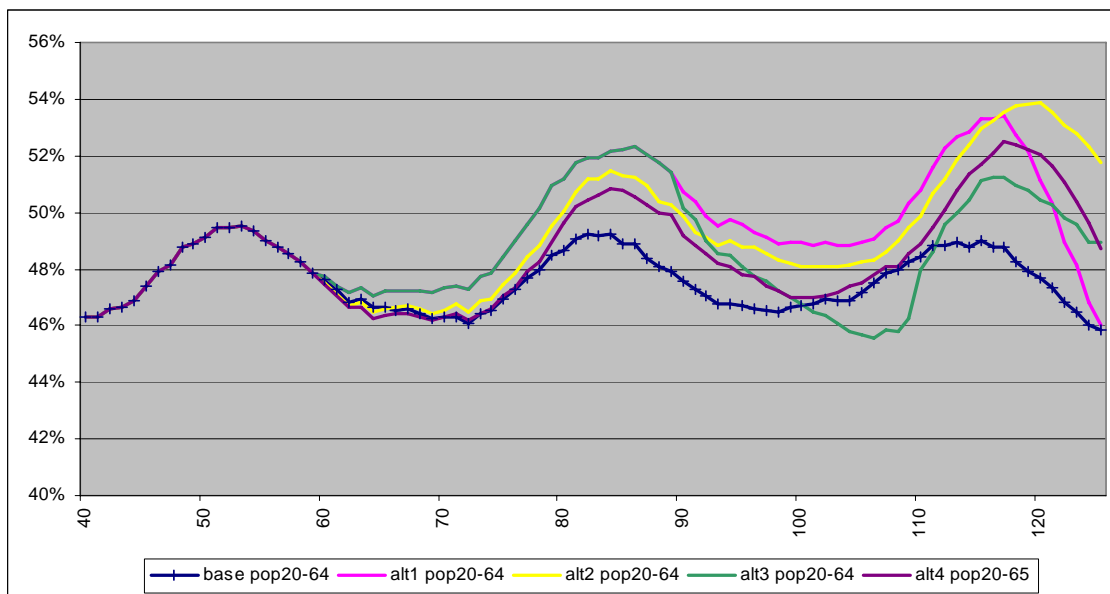


Figure 9 Share of working ages (20-64) relative total population

The population share in working age 20-64 increases in all alternative scenarios with about 3 percent in 20 years time (Figure 9). As can be seen below this gives rise to boosted GDP growth and lower taxes since less spending is needed in the educational system.

In all alternative scenarios parents delay their first birth compared to the base scenario. In the period before the child cost change occurs, mothers are about 29 years old when they have their first child, see upper part of Table 2. When the change is introduced in the model year 60 births are postponed in the first decade by two years for the first and third scenarios while the other scenarios two and four delay the downturn somewhat during the first decade. After two or more decades, however, births are postponed to about the same extent in all four alternative scenarios. This reflects that aspired income is more difficult to reach with the same number of children as before and this deters households from having the same number of children as their parents.

Table 2 Average age of mothers at first birth in the base senario, and the difference compared to the different alternative scenarios, by model time

<i>Model year</i>	<i>Base</i>	<i>Alt.1-Base</i>	<i>Alt.2-Base</i>	<i>Alt.3-Base</i>	<i>Alt.4-Base</i>
50-59	29.2	0.0	0.0	0.0	0.0
60-69	29.4	+2.0	+1.3	+2.0	+0.5
70-79	29.4	+2.5	+2.6	+2.5	+2.3
80-89	28.9	+3.2	+3.5	+3.2	+3.4
90-99	30.0	+3.7	+3.4	+3.5	+3.3
100-109	29.3	+4.5	+4.8	+4.5	+4.9
110-119	28.8	+5.3	+6.4	+6.7	+5.4

The human capital growth development, portrayed in Figure 10, suggests that the human capital stock is growing faster in all four alternative scenarios. This is a self-enforcing process. Once education has started to increase; the increase is sustained by a higher level of human capital in the economy making education more efficient.

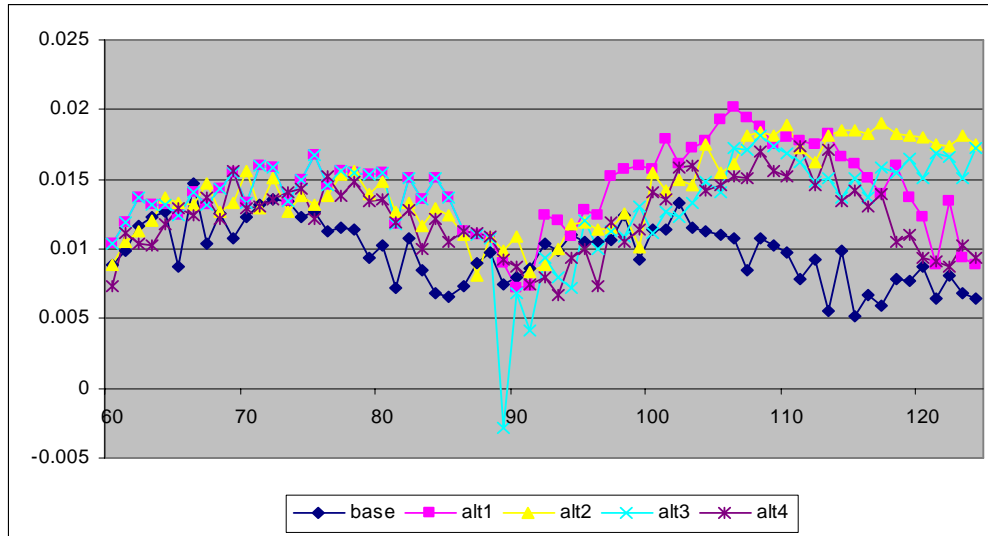


Figure 10 Human capital growth

Last, the growth of the economy in the scenarios 1-4, depicted as per capita GDP and shown in Figure 11, exhibits the expected pattern. The alternatives with lower fertility at first allow a much faster growth. As mentioned the relative cost mechanism is set off by the GDP growth. High GDP growth scenarios therefore creates a self-generating process of high child costs, further reducing fertility. The high per capita growth compared to the base scenario lasts in some scenarios almost 100 years (but one should note that GDP is not growing that long in levels, due to shrinking population). This raises the issue of how far into the future the altruism of the current generation will stretch. By abstaining from children they can improve their material well-being but only at the expense of as yet unborn generations in the future.

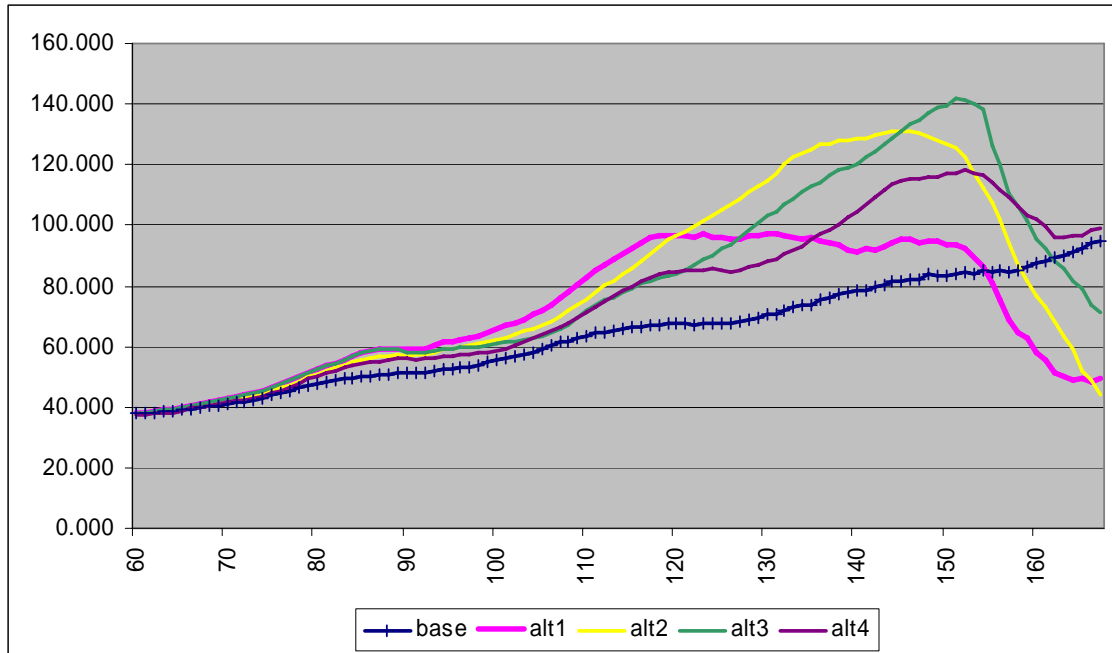


Figure 11 GDP per capita in 000nds

In Figure 12 depicts tax rates start increasing simultaneously as per capita growth declines. Tax ratios then quickly increase to a hundred percent of GDP. As expected this happens earlier in the “do nothing” scenario (alternative 1). The very large child benefit introduced in scenario 4, which at least temporarily seems to stop the drop in fertility, will imply a very high price of about 10 percent, or more, higher tax rate *in all future periods*. Implementing this policy would not be politically possible unless median voters have a very high degree of altruism towards future generations. In sum, however, high tax increases like these may in the long term not be fiscally sustainable anyway and therefore tax hikes to provide high child benefits cannot by themselves reverse the trap. Introducing saving and capital investment might change this conclusion (see Mason and Lee, 2007, about the second demographic dividend) but this is another paper to write.

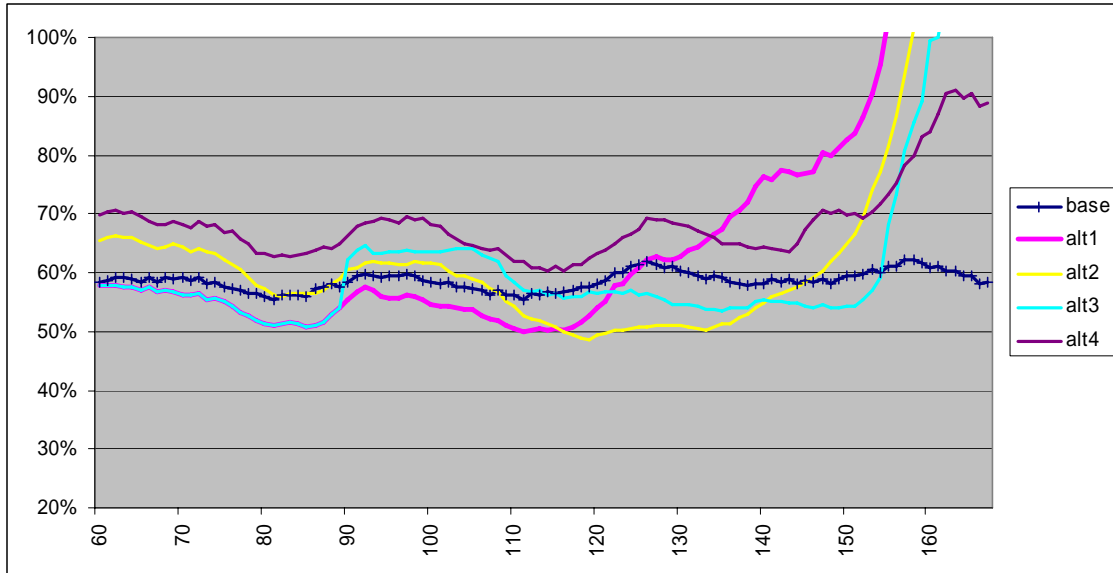


Figure 12 Tax revenues per GDP (“tax-ratio”)

Importance of social factors

How much of the fertility trap that we have created is a result of economic feedback from one’s parental home while growing up and how much is a result of the gradual formation of social norms? Some of the answers may be pulled out of Figure 13. This shows alternative scenario 5 where the social norm adaptation to fewer children has been shut down, but we have maintained the relative cost mechanism of having children just as in the other alternative scenarios 1-4.⁶

In contrast to alternatives 1-4 there is no clear drop in the population in this scenario, but instead a mild increase although smaller than the base scenario. Cohort fertility is a little lower than the base but still above levels of reproduction. This scenario bears a close resemblance to the base scenario also in other dimensions, like age dependency ratios, per capita GDP growth etc. Although a new child will be viewed as more expensive, this affects outcomes more like an idiosyncratic disturbance to adapt to. It seems thus that the social norms as we have modeled them have great impact on the development, reinforcing and forever letting small imbalances spread both through the population and over time.

⁶ Social norms work in two ways. First, the average number of children in the network influences the preference for children, and, second, the number of siblings to the spouses determines the initially preferred number of children. These two links are turned off in the fifth alternative scenario. See appendix for more exact details.

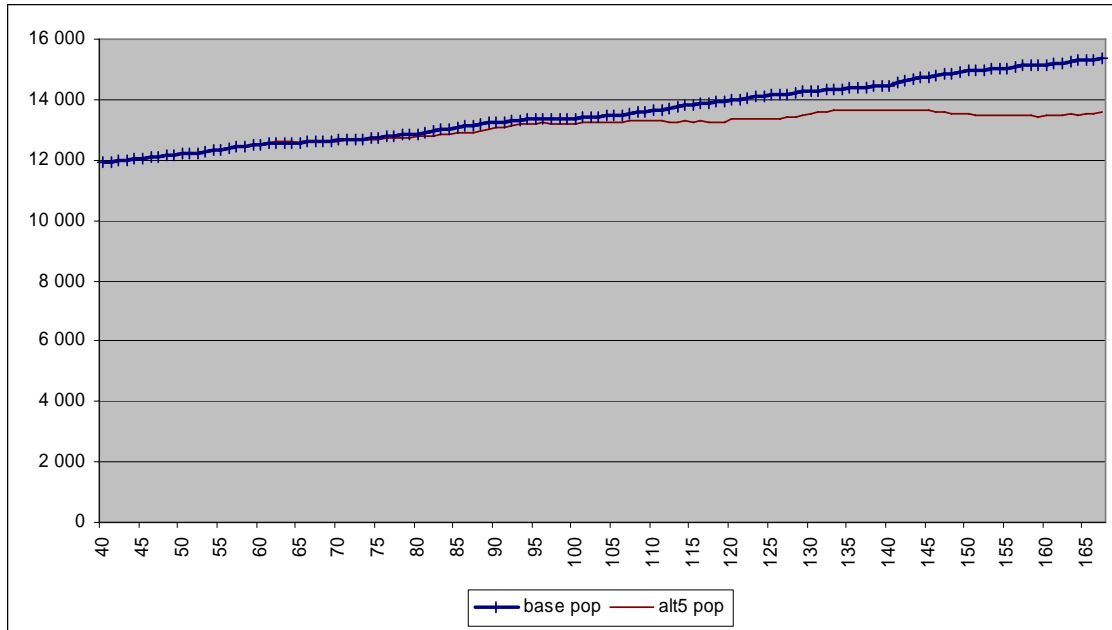


Figure 13 Population in base scenario and in an alternative without social norms influencing the fertility decision

Conclusions

Our models are capable of reproducing a fertility trap without crashing or demonstrate counter-intuitive behavior over a secular period of time. The alternative scenarios show that to get out of the trap requires very determined and persistent policy measures with high temporary growth costs, compared to doing nothing. Temporary in this context implies a period longer than the expected remaining life time of the currently active population. Without a fair degree of altruism such measures would not be politically feasible. Selfish individuals without consideration for their offspring far into the future would delay action until recovery is impossible because the tax base has grown too thin to support the necessary transfers. As we remove the social norm dynamics it turns out that this by itself prevents getting trapped; the economic mechanisms can adjust to regain balance in spite of the disturbances. It is the social norm dynamics which entrench new behavior in the population which is ultimately unsustainable.

Of course, our conclusions are only strictly valid for our virtual world. In the real world there are a lot more margins to adapt on. Our virtual world is, however, reasonably complex and still reproduces features of demographic dynamics in Sweden that we had not expected. We therefore conclude that the following hypotheses are worth further research.

1. Within an isolated system of intergenerational transfers where relative costs of children are increasing a low fertility trap can form.
2. When this process has been entrenched in social norms it is very costly to reverse the trend.

3. Unless voters are fairly altruistic this will not be politically feasible in a democracy.
4. A dissolution of the social norm mechanism seems the most important policy to pursue, for example by increasing social mobility.

Future empirical research should pay more attention to measuring the relative costs of children and whether variation in this variable can be causally tied to actual fertility behavior. In simulation it is important to test whether a savings mechanism would modify conclusions. There are a number of reasons why capital markets may circumvent the set of vicious circles we have set in motion in our virtual world. First, capital investment may to some extent substitute labor. Second, the “second demographic dividend” of net saving for own pensions may circumvent the aspired income mechanism by providing sufficient returns to ensure achieving aspired consumption levels without refraining from either education or children. Third, a capital market would allow public borrowing to invest in children without further diminishing the income disposable for consumption.

Acknowledgements

We greatly appreciate comments from Vegard Skirbekk and participants in The International Conference on the Economic Consequences of Low Fertility in St Gallen April 2008. We are very grateful to Matias Eklöf, Gustav Öberg and Elisa Baroni who have programmed the simulation model IFSIM and with great patience implemented the experiments this paper builds on. Funding from the Swedish Research Council is gratefully acknowledged.

References

- Aghion, P., Carolí, E., & García-Peñalosa, C. (1999). Inequality and Economic Growth: The Perspective of the New Growth Theories. *Journal of Economic Literature*, 37(4), 1615-1660.
- Allais, M. (1947) *Economie et Intérêt*. Paris: Imprimerie Nationale.
- Auerbach, A. J., Gokhale, J., & Kotlikoff, L. J. (1992). Generational Accounting: A New Approach to Understanding the Effects of Fiscal Policy on Saving. *Scandinavian Journal of Economics*, 94(2), 303-318.
- Becker, G. S., Murphy, K. M., & Tamura, R. (1990). Human Capital, Fertility, and Economic Growth. *Journal of Political Economy* 98(5), Part 2: The Problem of Development: A Conference of the Institute for the Study of Free Enterprise Systems, S12-S37.
- Billari, F.C.; Prskawetz, A.; Aparicio Diaz, B., & Fent, T. (2007). The “Wedding-Ring“: An agent-based marriage model based on social interaction. *Demographic Research*, 17(3), 59-82.
- Björklund, A. (2006). Does family policy affect fertility? *Journal of Population Economics*, 19(1), 3-24.
- Bloom, D. E., Canning, D., & Sevilla, J. (2003). The Demographic Dividend: A New Perspective on the Economic Consequences of Population Change. MR-1274, RAND Population Matters Monograph, Santa Monica.

- Caldwell, J., Caldwell, P., & McDonald, P. (2002). Policy responses to low fertility and its consequences: a global survey. *Journal of Population Research*, 19(1), 1-24.
- Davies, J. (2004). *Microsimulation, SGE and Macro Modelling for Transition and Developing Economies*. United Nations University / WIDER.
- Diamond, P. A. (1965). National Debt in a Neo-Classical Growth Model. *American Economic Review*, 55, 1126-1150.
- Easterlin, R. A. (1961). The American Baby Boom in Historical Perspective. *American Economic Review*, 51(5), 869-911.
- Ferrarini, T. (2003) *Parental Leave Institutions in Eighteen Post War Welfare States*. Swedish Institute for Social Research, Doctoral Dissertation Series no. 58.
- Flood, L., Klevmarcken, N. A., & Olovsson, P. (1996). *Household Market and Nonmarket Activities (HUS), Volumes III– VI*. Department of Economics, Uppsala University.
- Guo, Z., & Chen, W. (2007). Below replacement fertility in mainland China. In Zhao, Z., & Guo, F. (Eds.), *Transition and Challenge: China's Population at the Beginning of the 21st Century*. Oxford University Press.
- Gustafsson, S. (2001). Optimal age at motherhood. Theoretical and empirical considerations on postponement of maternity in Europe. *Journal of Population Economics*, 14(2), 225-247.
- Kelley, A. C., & Schmidt, R. M. (2005). Evolution of Recent Economic-Demographic Modeling: A Synthesis. *Journal of Population Economics*, 18, 275-300.
- Klevmarcken, N. A. (2002). Statistical inference in micro-simulation models: incorporating external information. *Mathematics and Computers in Simulation*, 59(1-3), 255-265, [http://dx.doi.org/10.1016/S0378-4754\(01\)00413-X](http://dx.doi.org/10.1016/S0378-4754(01)00413-X)
- Klevmarcken, N. A., & Olovsson, P. (1993). *Household Market and Nonmarket Activities, Procedures and Codes 1984–1991*. The Industrial Institute for Economic and Social Research, Stockholm.
- Krueger, A. O. (1968). Factor endowments and per capita income differences among countries. *Economic Journal*, 78, 641-659.
- Kotlikoff, L. J., Smetters, K., & Walliser, J. (2001). The Coming Generational Storm. *Computing in Economics and Finance 2001*, Society for Computational Economics 276.
- Lee, R. D., & Mason, A. (2004). *A Research Plan for the Macroeconomic Demography of Intergenerational Transfers*. National Transfer Accounts Working Paper No. 1, <http://www.schemearts.com/proj/nta/web/nta/show/Working%20Papers>.
- Lindh, T., & Malmberg, B. (2007). Demographically based global income forecasts up to the year 2050. *International Journal of Forecasting*, 23(4), 553-567. <http://dx.doi.org/10.1016/j.ijforecast.2007.07.005>
- Lutz, W., Skirbekk, V., & Testa, M. R. (2006). The low fertility trap hypothesis: Forces that may lead to further postponement and fewer births in Europe. *Vienna Yearbook of Population Research 2006*, Austrian Academy of Sciences, Vienna.
- Macunovich, D. J. (1998). Fertility and the Easterlin hypothesis: An assessment of the literature. *Journal of Population Economics*, 11, 53-111.
- Macunovich, D. J. (2002). *Birth Quake*. University of Chicago Press.
- Mason, A., & Lee, R. (2007). Transfers, Capital and Consumption over the Demographic Transition. In Clark, R; Mason, A. & Ogawa, N (Eds) *Population Aging*,

- Intergenerational Transfers and the Macroeconomy. (pp. 128-162). Cheltenham, UK: Edward Elgar Publishing Ltd.
- McDonald, P. (2005). Low fertility in Singapore: Causes, consequences and policies. Paper presented at the Forum on Population and Development in East Asia, Beijing, May 16-17, 2005.
- Richiardi, M., Leombruni, R., Sonnessa, M., & Saam, N. (2006). A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulation*, 9(1), <http://jasss.soc.surrey.ac.uk/9/1/15.html>.
- Samuelson, P. A. (1958). An Exact Consumption-Loan Model of Interest with or Without the Social Contrivance of Money, *Journal of Political Economy*, 66, 467-482.
- Schelling, T. C. (1969). Models of Segregation. *American Economic Review*, 49(2), 488-493.
- van de Kaa, D. J. (1987) Europe's Second Demographic transition, *Population Bulletin* 42(1), Washington, The Population Reference Bureau.

Appendix

Model overview

In this section we provide a more detailed description of how we have modeled the fertility decision and the network composition. For a detailed description of the overall model we refer to the IFSIM handbook which is available on request from the authors.

Figure 15 presents the flow chart with the main elements of the model. We start by creating an initial population from the Swedish micro data set HUS. This is a representative sample with about 3000 individuals which we scale to obtain out wanted number of individuals in the simulation. In the first period of the model we also introduce some initial aggregate variables from Statistics Sweden, e.g. mortality rates. Then we start the simulation by assigning each individual to a network group (see below for more details). Second we age the individuals by one year and then decide if they will die or not based on actual and predicted mortality rates from Statistics Sweden. If the individuals are still alive they will have the possibility to get a child according to the fertility module which we describe below. Individuals that still live with their parents may leave and create a household of their own, but only if they are above the age 16. The final individual decision within the demographic module is if they will match with a partner.

When individual characteristics regarding demographic variables have been established the simulation continues to determine individual status regarding education. We determine if they are enrolled in education, and if they are at which level. The education module also determines if an individual will start a certain education and what their highest completed education level is. Next every agent goes through the labor market module. At this stage this module does not contain any direct choice variable for the individuals. If they are not in education or at parental leave and below the exogenous retirement age they work full time. They are then either employed on the regular labor market receiving earnings based on their human capital level or they will randomly be assigned to become teachers in which case they would be paid a teacher salary. The wages for those with only basic education vis-à-vis those with higher diploma are determined according to their marginal productivity. Depending on the individual's status during the period, education, work, parental leave, etc., his human capital will increase or decrease.

When all the individuals within the system have obtained their characteristics it is possible to for the state to calculate the public cost in term of pensions, education cost, child allowance and parental leave cost. To be able to cover the expenses the state will choose the tax rate endogenously. Note that there is no direct labor supply effect from raising taxes in terms of reduced hours of work or early retirement. There is however a labor supply effect through the education system which affects the start of the working period.

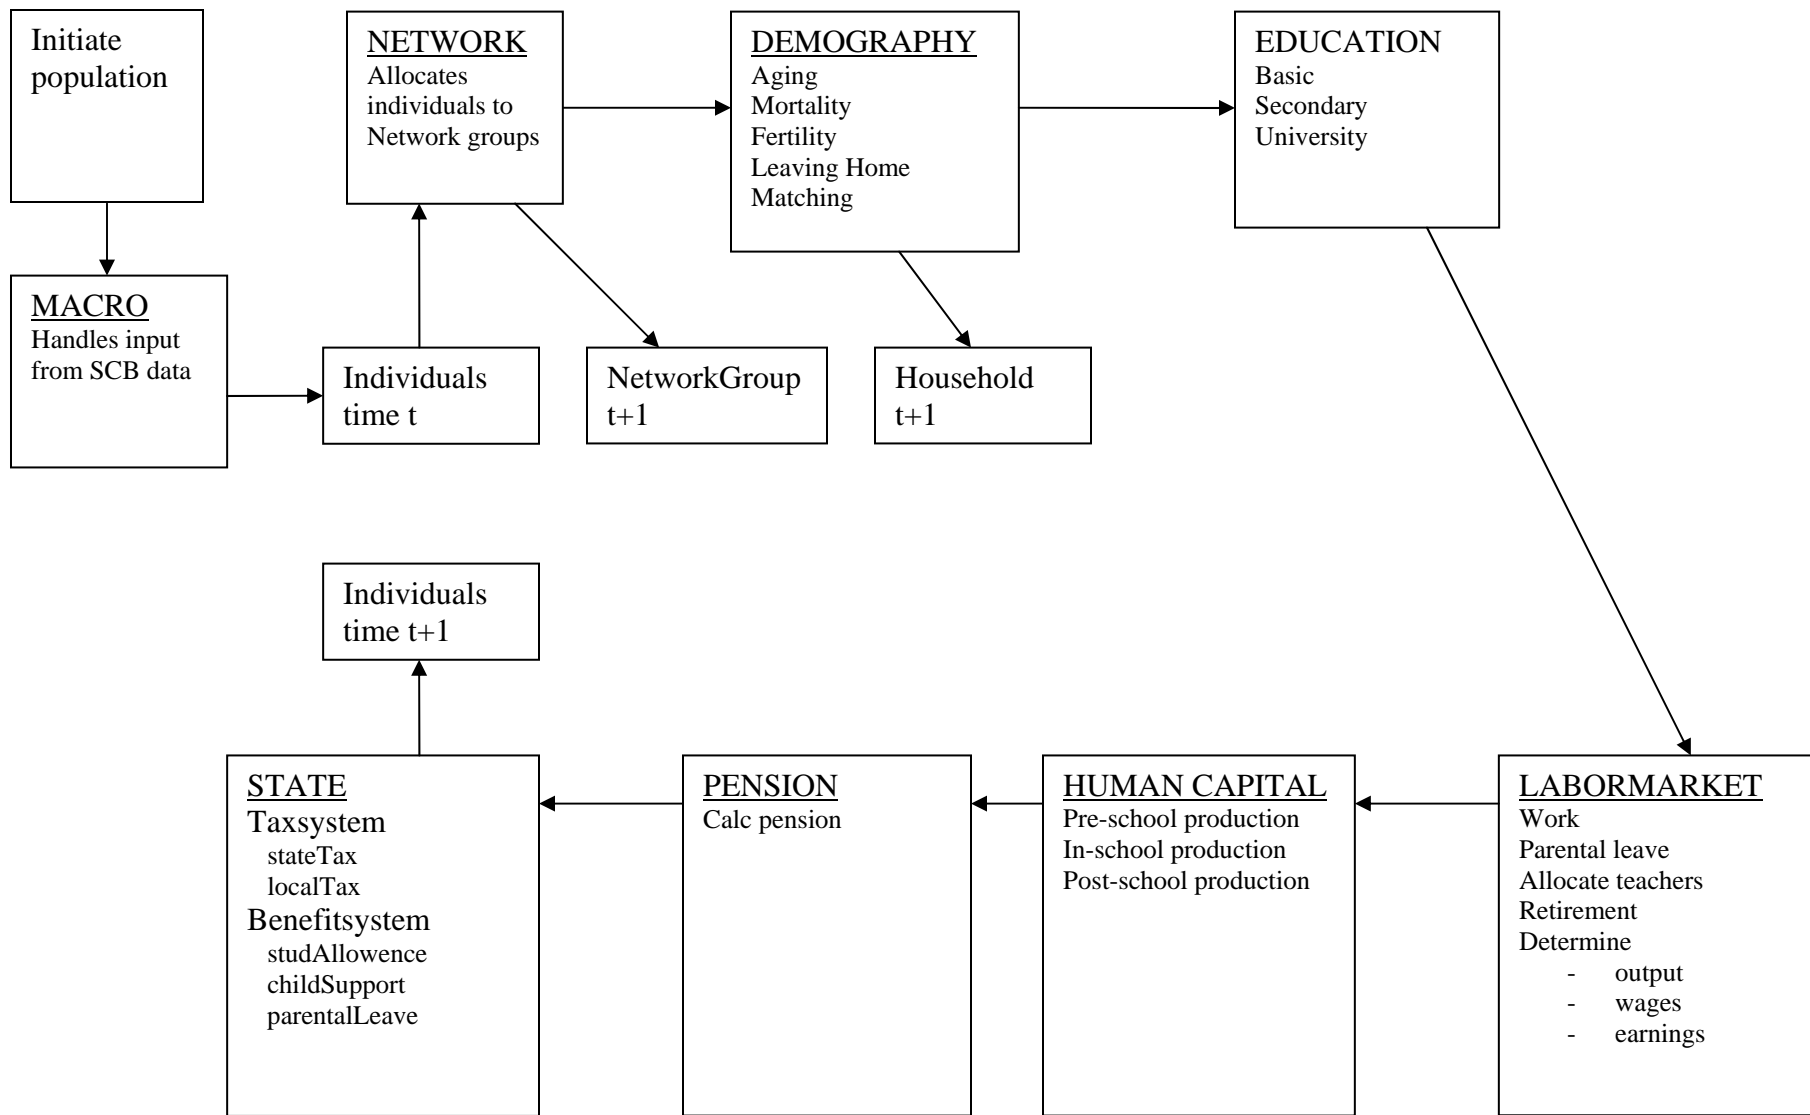


Figure 15: Flow chart for IFSIM.

Network

A distinguishing feature of ABMs is their ability to capture agents social interactions as these are supposed to influence individual decision making. The role of these social interactions is mostly to provide individuals with incomplete (as opposed to perfect) information which individuals in turn use to make decisions. This exchange is a process often described as social "learning", or even as social or peer "pressure" (Billari, Prskawetz, Diaz and Fent, 2006). For instance, the share of people married among one's friends might contribute positively to that individual's desire to get married herself, representing a conditioning or pressure to conform. The crucial idea behind modeling social interactions more generally is that this might indeed work together with economic incentives in explaining human behaviors, possibly affecting the size or even the direction that economic incentives might have otherwise. Our assumption therefore is that missing out social interaction from an analytical framework might bias the final results. In other parts of the model we use the decision of whether to go to gymnasium as an example of how to integrate social pressure into the more standard forward looking economic calculation of individuals when they decide whether to continue investing in their education or rather going into the labor market.

Every individual in our model is, from birth, member of a "social network" containing all those individuals to whom he or she is "close". We follow Billari, Prskawetz, Diaz and Fent (2006) in defining social "closeness" as a spatial area representing the individual's scope of interaction, by age group. Most specifically, agents are arranged along the surface of an imaginary cylinder, whose vertical length is broken into as many segments as there are age groups in the model (at present they are 8, from age 0 to Age_{Max}). Each age group therefore is allocated to an imaginary sub-cylinder whose height is the age interval for that group, and whose circumference is in turn sliced into a different number of networks (i.e. different age groups have different numbers of networks belonging to them). Each network is constructed as a segment on the circumference delimited by a corresponding angle. The model develops a procedure to then allocate each individual to a given network group within his or her own network space, by age group, and also to update his or her network in time, as the individual ages and moves between age groups and networks. A graphical representation of the network group organization is presented in the figure below.

This means that the individuals will migrate between network groups as the age, and two individuals that belongs to the same network group at one age, may belong to different groups at later ages. Furthermore, this implementation will allow for "spatial" migration as well as individuals could be allowed to change their "spatial" location, here measured as the angle on the circumference, over time.

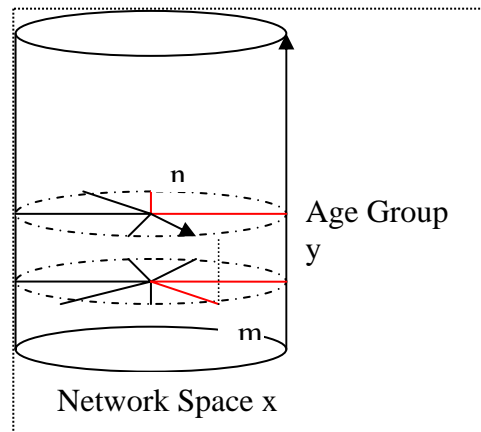


Figure 16: Representation of network groups

Any newborn inherits a "network angle" which is the average of her parents' plus a random term. A network angle is a technical construct which allows the individual to be placed, given his age, on a specific segment of the circumference throughout her life. At any point in time, his or her individual angle is compared against the angle determining where each network begins and end (by age group), and placed within the network segment corresponding to his or her angle. In the same network the individual will therefore meet a sub-sample of other people in the same age group who happen to have a similar angle (both randomly and due to parents' characteristics). In time, as the individual grows older and jumps between age groups, the composition of his or her network will change since different age groups are characterized by different network angles, hence individuals are shifted not only vertically but also "horizontally", so to say, depending on their own angle relative to the new age-group network configuration.

To sum it up, the individual's location within the cylinder space is determined by 2 coordinates: (i) their age group. This defines the location within the vertical Y space, or in other words the sub-cylinder to which the individual belongs (ii) the angle of the circumference within which the individual's network falls (corresponding to the interval on the circumference occupied by that network). This defines the horizontal x-coordinate.

By looping over all age groups and angle group, networks are thus populated. The size of networks will vary in time while the characteristics will remain relatively stable at least in terms of age composition. Each network group is modeled as a Java object capable to iterate over its members and extract a number of summary statistics such as averages by group (i.e. proportions).

Fertility

There are several variables that affect the fertility outcome. First we have fecundity which is more or less beyond individual control. We roughly capture this biological restriction by putting an upper limit at which females in our model can give birth. This upper limit is set to age 40, above which no female can give birth. We also have a lower

starting age set to 20 under which it is not possible to give birth. This lower limit has not that much to do with fecundity but we have chosen it since few give birth at such young ages.

Beside the biological restriction to giving birth we model a fertility function that allows us to capture the main elements of the low fertility hypothesis. The LFTH has a social dimension and an economic dimension. By the social dimension we refer to how our desire to acquire children is influenced by the number of children around us. We first start by assuming that the norm of how many children one wants to acquire is set during youth. We call this the wanted number of children. Every individual has a wanted number of children that is determined by the number of siblings that he/she had when young. For a couple we use both the males and females wanted number of children to construct the couples wanted number of children. This is the average number of the couple weighted with a random number. The wanted number is something that the couple strives to obtain but it is not sure that they will reach this number due their economic outcome and the social influence from their network group. Given that the female is in fertile age and that the number of wanted children is higher than the actual number she will give birth if the following conditions hold:

$$\frac{DISP}{\sqrt{n+r}} \geq \text{median income},$$

$$\text{SocialFactor} * PROJ^{10} > ASPC,$$

where $DISP$ is the disposable household income, n is the number of individuals in the household and r is the cost of a new child. We use the square root to capture the economics of scale of large families. This equivalent income must exceed the median income which implies that today's economic conditions are very important.

The second restriction we have is that the social norm in combination with expected future income must exceed the aspired consumption level. We define the social factor according to:

$$\text{SocialFactor} = \left(\frac{e^{N_{\text{kids}} N_{\text{members}}}}{1 + e^{N_{\text{kids}} N_{\text{members}}}} \right)^{\varphi}$$

where N_{kids} is the number of kids belonging to the N_{members} members in the network group. φ is a parameter that controls the strength of the social pressure currently set to 0.92. If many individuals within ones network have children this would positively affect the fertility decision of a couple.

The economic dimension consists of determining if the couple can afford a new child or not in the future. We follow the LFTH and state that a couple aspires a certain consumption level. They will not acquire a child unless they can reach this aspired consumption level. What remains is to define this aspired consumption level. As

postulated by Easterlin this aspired consumption level is a norm that is formed during youth and which has the consumption level of ones home of origin as the reference point. We model this aspired consumption, $ASPC$, according to:

$$ASPC = \theta \bar{C},$$

where $\bar{C} = \left(\frac{DISP}{\sqrt{n}} \right)$ is the average equivalent disposable income today and

$$\theta = \gamma \left(C_{at10} / \bar{C}_{at10} \right) + (1 - \gamma) \left(C_{at10} / PROJ C^{10} \right).$$

C is the equivalent disposable income for every individual in the household. The subscript $at10$ indicates that it is when the individual was at the age of 10, which captures that aspirations are set during youth. We also apply a weighting factor, $\gamma \in (0,1)$, for two different reference points for the aspired consumption. The first term states that the position in the consumption distribution at the age of 10 affects the aspired consumption level. The idea is that ones children should not obtain a worse position in the consumption distribution. The second factor captures the idea that parents do not want their children to have less consumption then what they had when young adjusted for economic growth. Since the reference point was set when the parents where of age 10 it is natural to compare the new child's consumption level at the age of 10. For this reason they project expected future income according to:

$$PROJ C^{10} = \frac{DISP^{10}}{\sqrt{n+r}},$$

Where $DISP^{10}$ is the disposable income ten years from today *if they choose to have one additional child*. In the base scenario, r is equal to unity, but varies between 1 and 3 in alternative scenarios, depending on economic growth in society. The household's disposable income in ten years is estimated from the sample of individuals whom today have ten more years of labor market experience. We correct for the fact that these individuals live in households with different age composition of the children (since they will, for instance, be on average 10 years older and thus not receive the same amount in child benefits which changes the household's disposable income). Let $Y(X, E)$ be the net labor income after taxes for an adult, who has X years of experience after diploma year from education level E (=basic, high school, university). This is net of taxes, and does not include the child benefit b . The child benefit is introduced in some of the alternative scenarios as explained in the text. It is given to all households with children 0-17 years of age. So at each year all household receives $n_{0-17}b$, where n_{0-17} is # kids 0-17 in the household today.

Thus, for each model year and education level E , we estimate (by ordinary least square, OLS) $Y(X, E)$ as a quadratic function in X , for each of the (potential) parents. Then the predicted value of someone with 10 more years of experience with the relevant

education level is $\hat{Y}(X + 10, E) = \hat{\alpha} + \hat{\beta}_1(X + 10) + \hat{\beta}_2(X + 10)^2$, where $\hat{\alpha}$, $\hat{\beta}_1$, and $\hat{\beta}_2$ are OLS estimates.⁷

The projected disposable income is the household's disposable income in ten years with one additional child. This is then:

$$DISP^{10} = \hat{Y}_m(X_m + 10, E_m) + \hat{Y}_f(X_f + 10, E_f) + (n_{0-7} + 1)b,$$

where m =mother, f =father, and n_{0-7} is the number of children 0-7 in the household today.

⁷ Those who are enrolled in education and will reach the education level E within 10 years will have a negative value of X corresponding to years until graduation.