Modelling the Easterlin Effect Using Agent-Based Simulation

Jason Hilton
Centre for Population Change
University of Southampton

Motivating an Agent-Based Model of Intergenerational Ferility

This paper aims to examine the plausibility of an existing theory which relates cycles in macro-level fertility to an individual-level desire to delay childbearing until a level of well-being commensurate with that achieved by one's parents [Easterlin, 1987]. This phenomenon (dubbed the 'Easterlin effect' [Pampel and Peters, 1995]) does not occupy the minds of many fertility researchers at present, because it is of mostly historical interest; it is deemed to have operated in a period in the early 20th century to the late 1970s. However, it is an interesting target for a simulation study because it is a simple theory that explicitly describes links between the micro- and macro-levels in a demographic process, and because it has clear implications which can be tested (the generation of cycles in fertility). Furthermore, the 'relative well-being' mechanism through which the childbearing decision relies on individual relationships between agents - specifically, those between parents and children.

Modelling this mechanism directly in an Agent-Based Model, although not without its own challenges and shortcomings, avoids some of the specification and endogeneity problems found in more traditional empirical approaches to the problem, highlighted in review by Macunovich [1998] and Waldorf and Byun [2005], while recognising micro-foundations not present in the mathematical models of Lee [1974], Wachter [1991] and others. However attempting to model a complicated social phenomenon of this nature often requires the inclusion of several interconnected processes, each of which is governed by its own set of parameters, and each of which is a potential source of error and uncertainty. Thus, the use of statistical emulator techniques is important for allowing a systematic analysis and calibration of the simulation described in the subsequent sections.

Description and Theory

The Easterlin effect purports to explain the existence of distinctive wave-like patterns in the time series of births in the United States. It is associated with the the work of Richard Easterlin, who developed his theory in a series of articles and books from the late 1950s onwards [e.g. Easterlin, 1962, 1966, 1975, 1987]. These patterns are different from the expected generational cycles, whereby large cohorts have relatively more children than smaller cohorts simply because they include more prospective parents, and thus 'echoes' of larger cohorts are seen in population age structures at intervals of approximately a generation's length.

Instead, Easterlin is describing waves of double this period, whereby a child's fertility is likely to be negatively correlated with that of his parents' generation, but positively correlated with his grandparents'. Examining Figure Figure 1, net of the upward trend, broad arcs are visible in the plotted time-series of births, until some flattening off after 1980.

Easterlin posits that generational cycles such as those described are the result of a link between birth cohort size and later fertility, a link that is caused by reduced opportunities in many spheres for those in larger cohorts, and thus an increase in the probability that those in such cohorts will not feel well-off enough to start a family.

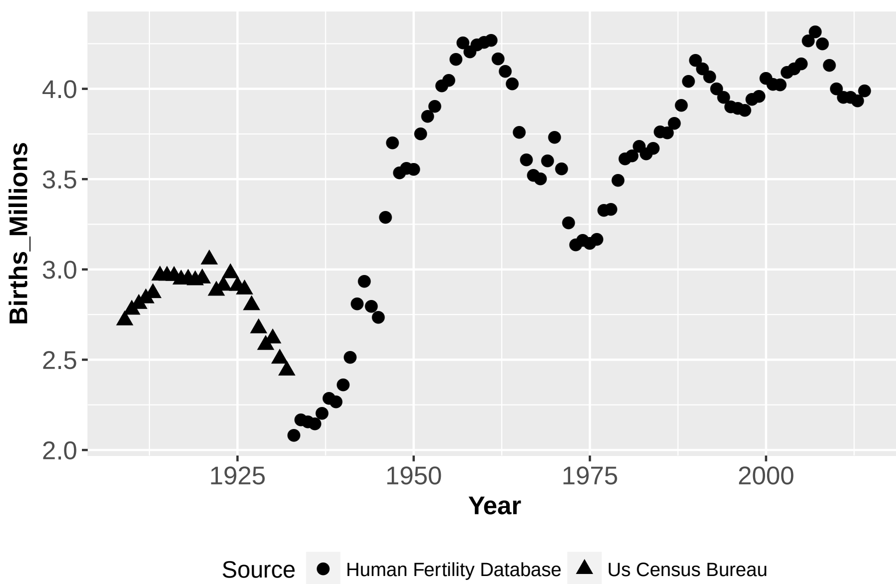

Figure 1: Time Series of US Births since 1909. Source: US Census Bureau, 1975; Human Fertility Database, 2015

To breakdown his thesis further, it is possible to identify four specific assumptions upon which it rests [Easterlin, 1987].

1. Individuals wish to start a family only once they have reached a level of well-being that satisfies their aspirations.
2. These aspirations are set relative to some reference group, rather than in terms of some absolute standard of well-being.
3. The relevant reference group against which individuals set their aspiration is their parents.
4. Individual well-being is negatively correlated with cohort size, in particular due to increased competition in the labour market.

This last statement contains a subsidiary assumption about the labour market, namely, that older and younger workers are imperfectly substitutable, or else relative cohort size would not effect the success of young workers [Macunovich, 1998]. The fact that older workers are likely to have obtained greater experience in their chosen career (equivalently - developed higher levels of job- or trade-specific human capital) goes some way to laying the ground for such an assumption; often, higher level jobs require experience to be performed effectively.

Easterlin also hypothesises that competition and crowding in two other areas of life may lead to poorer outcomes for those in large cohorts [Easterlin, 1987]. Firstly, larger families lead to greater divisions of parental resources, both monetary and in terms of time and attention. Secondly, crowding in educational establishments may also be problematic; if school capacity building and teacher training lags behind demand, then it is likely that smaller cohorts will be at an educational advantage. However, this paper focuses only on intra-cohort competition as it manifests itself through the labour market. Easterlin's hypothesis was somewhat heterodox in economic theory at the time because, rather than taking consumption preferences as given, it considered them as malleable and something to be explained [Easterlin, 2004]. Preferences are described as determined through a process of socialisation - being exposed to certain levels of consumption during later teen-hood make one desire to attain at least this level of well-being.

The combination of these factors, then, is hypothesised to results in cycles of two generations in length, as small cohorts have little job market competitions, and therefore outperform their parents' generation, and therefore give birth to relatively more children. The resulting larger cohorts are then less successful due to the increased competition for jobs and resources they face, and therefore are condemned to lower average fertility, and so on [Easterlin, 1987].

The relative lack of immigration during the period in question due to strict immigration laws is also described as providing the conditions necessary for such cycles to develop [Easterlin, 1978]. If immigration was subject to few restrictions, labour shortages due to smaller cohorts would not lead to a bidding-up of wages, but instead to a greater inflow of immigrants to meet the demand. Similarly, an relatively large cohort might lead to a cessation of such flows, leading to little change in the experience of labour market conditions between cohorts, regardless of size.

Cohort sizes are also described as having implications for female labour force participation, albeit in the opposite direction from those that held for men.The change in sign is a reflection of the reality of the more detached role of women in the labour market in the mid-twentieth century [Easterlin, 1978], largely due to the limited opportunities available to them as a result of discrimination and oppressive social norms. Older and younger women were supposed more substitutable as labour sources due to the 'non-career' roles they were typically obliged to take. Thus, scarcities in young male labour may lead to better employment prospects for them, but may lead to a withdrawal from the labour market of their spouses, as they were expected to devote themselves to motherhood, and an increase in employment among older women, who may be incentivised by higher wages to fill the gaps. Conversely, poorer prospects for the spouses of younger female cohorts may require them to work more - particularly as they are less likely to start families, with knock on effects for the demand of older female age groups.

Empirical Evidence in the literature

In part due to its apparent success in describing the American fertility boom and bust, Easterlin's theory has received a lot of attention in the demographic literature. Broadly speaking, this work can be split into two parts; studies that focus on the macro-level, looking at demographic and economic time-series and mathematical models of the process, and those which centre on the micro-level, and use survey data to investigate the relationships between parent's wealth, income, cohort size and fertility. Pampel and Peters [1995] and Macunovich [1998] provide comprehensive reviews of such studies.

The Easterlin Effect at the Macro-Level.

Evidence for the Easterlin hypothesis is mixed. At the macro-level, in terms of empirical data, it is clear that since the early 1980s the observed cycles have vanished, as can be seen in Figure Figure 1. A number of reasons for this breakdown have been suggested. Increases in immigration and unemployment together with less secure working arrangements, more temporary work and increased female labour force participation are just some of these [Pampel and Peters, 1995]. Indeed, Easterlin's original analyses (e.g. 1962, 1967, 1978) were very much concerned with the interplay of cohort size, immigration, and labour force conditions in different sub-groups in determining fertility, due to his initial association of fertility cycles with immigration-linked Kuznets waves in economic activity.

Despite this later failure of the theory's predictions, the relationship between cohort size and later fertility appears to hold strong for much of the twentieth century - fluctuations in fertility from about 1900-1970 follow relatively closely those of lagged cohort size [Wachter, 1991]. Using birth-order life-tables constructed by the Human Fertility Database Project, it is possible to examine how fertility rates differ by birth order as the swing from high fertility to low fertility progresses between 1963 and 1976 (Figure 2). Birth-order life-tables describe the rate at which those with n children transition to having $n+1$ children, based on data for single period. Those with n children are described in demographic terminology as being at *parity* n. Note that rates at all parities decline over the period, so that if relative income is responsible for the sharp downturn, it is likely to act on all decisions to have a child, not just the initial transition to parenthood.

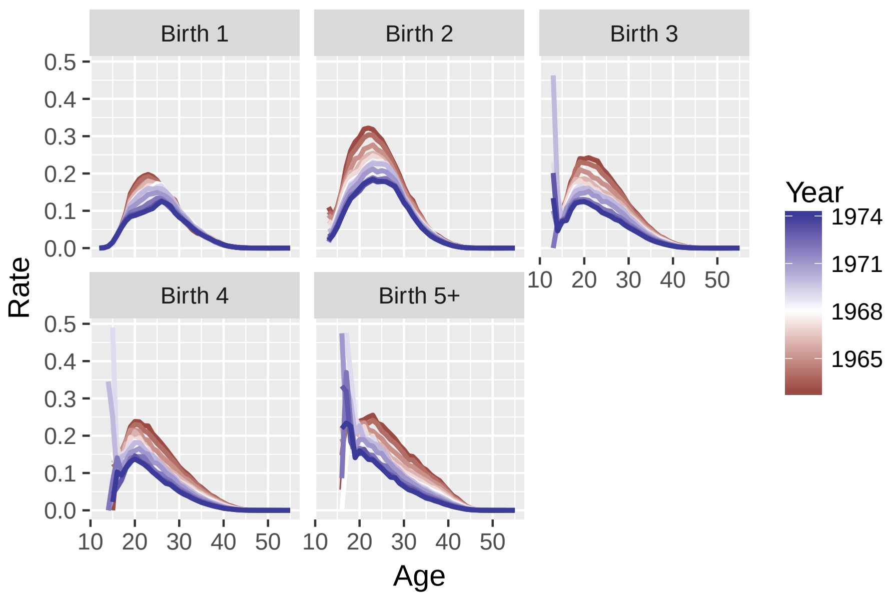

Figure 2: US Birth Rates by Birth Order. Source: Human Fertility Database, 2015

A considerable amount of work has been done creating macro-level mathematical models of the Easterlin process, most notably by Lee [1974], Frauenthal and Swick [1983], Wachter and Lee [1989] and Wachter [1991]. These examine formal, dynamical systems models where fertility rates are damped by larger cohorts and similarly boosted in the presence of smaller ones. Thus, fertility is affected to various extents by the past course of births, allowing deviations from the equilibrium path [Lee, 1974, Wachter and Lee, 1989, Wachter, 1991]. In general, these models involve formulations of the type shown in Equation 1.

$$B(t) = \int l_s m_s B(t-s) M \left[\frac{B(t-s)}{e^{r(t-s)}} \right] ds \tag{1}$$

where $B(t)$ denotes the number of births at time t, l the survivorship ratio, m a fixed component of fertility, and r the equilibrium growth rate. The significant part of the model is the expression $M[.]$, which is a function of the sizes of past cohorts $B(t-s)$ relative to the equilibrium trend $e^{r(t-s)}$. Various forms of this function have been analysed to see if the resulting system can sustain limit cycles, and whether the parameters on $M[.]$ are similar to those estimated from US data [Lee, 1974, Wachter and Lee, 1989, Frauenthal and Swick, 1983]. Wachter [1991] analyses a family of models of this nature, with fertility depending on either the whole of the labour force or on various subsections or cohorts. In general, he suggests that the existence of self-generating Easterlin cycles is possible, but that it is quite unlikely that the macro-forms specified are responsible for the patterns observed in the US [Wachter, 1991].

A study by Waldorf and Byun [2005] reviewed macro-level studies in a systematic manner using the tools of meta-analysis. The findings were mixed, with supportive results being more likely in the US, and highly dependent on the control variables used and the specific functional forms used. Interestingly, they suggest that (publication) biases towards negative findings (disproving the thesis) were more likely the more recently the study was published, while for older studies, the opposite bias was suggested. Furthermore, the authors found that the use of income as a control or explanatory factor substantially changed the direction of results, and highlighted endogeneity concerns in the use of this variable.

The studies discussed above utilise purely macro variables and do not discuss the underlying mechanism relating to fertility damping. Furthermore, they generally assume structural heterogeneity (that is, the parameters remain constant) across time.

The Easterlin Effect at the Micro-Level.

A large number of studies based on survey data have been carried out examining whether the Easterlin effect holds at the micro-level. Macunovich [1998] provides a review of these studies (together with some macro-analyses), and suggests that the evidence is mixed, and many studies do not find an association between relative income and fertility, particularly outside of the U.S. However, she believes that these negative findings are due in part to data problems and specification errors in the models used to test the theory. In particular, measurement of past parental wealth or income is often difficult, and where it is present as a data-set variable it is often only available for one parent. Furthermore, the assumption that a simple threshold, defined by material well-being in teen-hood, will suffice as the operationalisation of the theory is also critiqued by Macunovich; instead, some function of the past well-being may also be consistent with the theory.

More recently, research has been conducted into the possibility that happiness and satisfaction is often a prerequisite for family formation, Parr [2010], for instance found that fertility is related to prior satisfaction with life, although he does not discuss the cohort effect in determining such satisfaction. Teitelbaum and Winter [2013] describe a different perspective somewhat commensurate with Easterlin's thesis; that declines in fertility can be understood as a risk management strategy. Thus, even if not directly affected by the detrimental effects of being born into a large cohort, perceptions of the risk of suffering unemployment or stagnant wages may still relate cohort size to lower fertility indirectly.

Overall, there is no clear picture in the literature as to whether Easterlin cycles were indeed behind the boom and busts in fertility observed in the US between about 1900 and 1980, in part because of the difficulty in generalising from such a short-lived phenomenon, data problems, difficulties in specifying the model, and so on. This paper approaches the problem from the opposite direction, as will be discussed presently.

An Agent-Based Approach

The focus of this paper is on attempting to 'grow' Easterlin cycles from the bottom up [Epstein and Axtell, 1996]. That is, by specifying the micro-level behaviour of agents in a simulation in line with the mechanisms suggested by Easterlin, we attempt to recreate the observed cyclical patterns. If we can succeed in doing so under reasonable assumptions, we can assume that Easterlin's micro-level specification is plausible. Thus, the research question is as follows:
1. What behavioural rules and micro-level conditions are sufficient to generate Easterlin-like waves in fertility?

Methodology

In order to examine the interplay between the population and individual level inherent in Easterlin's theory, and to capture the influence of population heterogeneity and intergenerational links on the relative well-being thesis, a discrete time agent-based simulation was built. This allows us to formalise the particular mechanism suggested by Easterlin, maintaining the specific links between generations that result in the cycles observed in the data.

The model simulates the life histories of agents as they are born, age, find jobs, start a family, and eventually die. The model is relatively simple; the idea is not to capture every element of social life with absolute fidelity, but rather to distil these elements to the essences which are required to study the question in hand. However, an attempt is made to include population heterogeneity, particularly in earnings, and explicit relationships between agents as integral parts of the model, in a way that is more difficult in other modelling paradigms.

The logic of the overall approach is as follows. The simulation represents a idealised abstraction of the hypothesis given by Easterlin; if it is successful in replicating the waves in fertility of similar period to those seen in the twentieth century, then we can consider the theory plausible. However, the process of building an Agent-Based Model is far from trivial. Many assumptions about specific encodings of individual behaviour must be made, and additionally, many of these encodings will involve parameters for which we do not know the true empirical value. This means that the simulation must be calibrated in an attempt to match the empirical results. This requires the use of statistical techniques, as finding distributions of well-fitting parameters in high dimensional spaces with non-linear response profiles, as is common in Agent-Based Models, can be difficult. Gaussian Process emulators are thus used in order to identify plausible parameter ranges given empirical observations, in line with work by Vernon et al. [2010] and Boukouvalas et al. [2014]. The next section discusses the architecture and specification of the model in detail.

Overview of the Simulation Model

The extensive nature of the mechanisms Easterlin described as driving his cycles, incorporating labour market participation, wage setting, fertility, partnership and other elements besides, mean that a complete description in a single model is probably undesirable, at least in a project undertaken by a single programmer/analyst. The decision is thus made to treat certain elements of the model as exogenous, and in particular, the simulation is so engineered so that wages and unemployment are assumed to respond to cohort size exogenously, and in the way Easterlin described. This allows the focus to remain on the key element, that of fertility choice.

Simulation Structure

A simulation was built to reflect the key assumptions of Easterlin's theory, namely, that fertility is related to relative income, which is in turn related to cohort size. This sub-section provides a high-level overview of the simulation structure, while the next section provides a detailed description of each element of the simulation. The code is freely available as a `git` repository at https://bitbucket.org/jhilton/easterlinphd, and requires only `python` together with some easily downloadable libraries to run.

Model Specification

A description of the key elements of the simulation is now offered, together with how they relate to the ideas in Easterlin's thesis. The simulation initialisation process is also described, together with a summary of how the model executes.

The Agent

The basis of the model is the agent. Males and Females are represented by different sub-classes in the code, although they share many of the same methods. Instances of the agent class keep track of the key state variables in the model: age, the identity of an agent's immediate family and an agent's individual 'skill' level, reflecting their ability to earn a wage in the market. This is drawn from a normal distribution centered around 0.5, and truncated close to zero and at one:

The class also contains methods relating to some of the exogenous demographic processes not of central interest to the research, including mortality and partnership. Mortality is described by a Gompertz function, as displayed in Equation 2, while mortality below aged 30 is assumed zero for the sake of simplicity, as it is expected that this will not make a noticeable difference to the outcome of the model. Differences in mortality by sex are ignored in the simulation, and while the Gompertz model parameters are configurable, they are kept fixed for the experiments detailed below.

$$\begin{aligned} m(x) = a_m \exp\left(b_m x\right) \\ \forall x : x > 30 \end{aligned} \tag{2}$$

where a_m and b_m are further parameters describing the slope and intercept of the log mortality function.

Partnership is also considered exogenous, but is rather more complicated in its implementation. Following Zinn and Himmelspach [2009], agents enter the marriage market according to an age-specific hazard modelled by a double exponential function as shown in Equation 3. Once in the marriage market, agents undergo a matching process, whereby females in the market, by order of entry to the market, are able to choose the most compatible male partner by age and skill, determined by a Euclidean distance measure. The ideal degree of age difference between partners is a parameter in the model, but is held constant at 3 years for the purposes of this work. The age-specific schedule is also parametrisable, but is set up to ensure marriage is relatively early, in line with the norm in the early twentieth century, and to ensure that the major determinant of fertility is the relative income mechanism described by Easterlin, and not the rate of entry into marriage (although this could also depend on relative income and cohort size to some extent, as better-off males might be more suitable mates and more inclined to set up a family [Easterlin, 1987]).

$$n(x) = a \exp(-\alpha_n(x-\mu) - \exp(-\lambda_n * (x-\mu))) \tag{3}$$

Population

The population class exists mainly as holder for the collection of agents, and coordinates their actions and interactions. Agents are held as a `python` list object within the class, and this list is iterated through allowing agents to undertake their relevant yearly actions. The order of iteration is randomised every time-step so that no particular agent always 'acts' first.

Labour Market

More complex is the labour market. Earlier experiments with this simulation attempted to include labour market processes as an endogenous part of the model [see also Fagiolo et al., 2004], with firms and agents setting desired prices for labour based on their knowledge about supply conditions, which could be learnt from interactions with other agents. However, while this is an interesting avenue for research in its own right, it does not form the central focus of this investigation, which is more concerned with fertility decisions.

The labour market element of the simulation, then, is intended to reflect the ideal conditions under which Easterlin-like cycles might be expected to flourish. In particular, wages and employment opportunities are set up so that those in smaller cohorts are more likely to get jobs, and the jobs they do get will be better paid. Note that it is by no means certain that this is how the labour market does behave - while the supply of labour is expected generally to affect labour, other factors, such as firms' capital investment responses to changes in labour supply and the possibility of increased immigration, might be expected to mitigate this effect. Furthermore, exogenous processes such as technological advances and longer-term changes to the relative market power of capital and labour also may interfere with this relationship [Cahuc et al., 2014].

To turn to the specific implementation details, the number of jobs of the economy is held to be proportional to the population, with a weighting towards those of working age in line with the consumption patterns identified by Lee et al. [2011]. An agent's capability to produce in a given job is defined by a fixed productivity function, which varies according to the amount of experience an agent has in the labour market, the agent's predetermined 'skill' level, and the 'difficulty' of the job. Experience (defined as the number of past years of employment) is assumed to affect productivity in line with the Mincer model of lifetime earnings [Mincer, 1974, Cahuc et al., 2014]. More specifically, productivity increases with experience at a decreasing rate, before declining for higher values, reflecting the acquisition of human capital throughout an individual's life while allowing depreciation of same as retirement approaches (ibid).

Skill and difficulty are related to productivity in such a way that higher skilled agents have a productivity advantage when undertaking more difficult tasks [Cahuc et al., 2014, section 10.2.2]. The specific functional form of productivity is given in Equation 4, with s, d and e representing skill difficulty and experience respectively. The final choice of function is somewhat arbitrary, but has a basis in economic literature and results in a realistic log-normal distribution of wages for many populations and parameter choices. Including graduation in jobs and earnings within the labour market allows the potential to examine how distributional aspects might affect the workings of the mechanism identified by Easterlin.

$$p(e, s, d) = \beta d \exp(\alpha(s - d) + s + \gamma e - \delta e^2) \tag{4}$$

Realised wages in the simulation are deterministically related to the size of a particular cohort by means of an adjustment to this productivity function. The relative size of each working-age birth cohort is calculated, and a Gaussian kernel centred on an agent's age is used to take a weighted average of nearby age groups to arrive at a measure of relative cohort size. The final wage (given in Equation 5) is then realised as the product of the exponential of this weighted cohort size f and the productivity p given by Equation 4. In addition the elements l and m allow for the possibility of linear and exponential growth in time t, reflecting the effect of economic growth. Thus, the effect of supply of labour on wages is treated as given and exogenous in this model.

$$w(p, f) = lt + \exp(\zeta f + mt)p \tag{5}$$

A number of factors are ignored in the model. Specific relationships between earnings, consumption, savings, and demand for labour are at present ignored, for the sake of focusing on the demographic rather than the economic elements of the phenomenon.

Employment

The employment class controls the labour market behaviour of the agents themselves. Agents begin to apply for jobs when they pass the age of 16, and retire from the market at age 65. Agents apply at random to vacant posts, and are offered jobs if they are the best candidate, assessed according to their contribution. Agents accept the job if it is the best offer they receive. The number of applications is drawn randomly each turn and depends upon an agent's employment status; those without jobs apply more widely. The mean number of applications made while employed or unemployed are tunable parameters in the model. A certain amount of 'churn' is also introduced, meaning that a small number of jobs are created and destroyed every time step, representing random exogenous shocks. At present, the simulation is run with only male breadwinners in an attempt to reflect the reality of the period, although female labour market participation is possible within the simulation, and increases commensurate with those occurring in the later twentieth century could be modelled in future work.

Fertility

The fertility class is central to the simulation behaviour. The modular nature of the simulation allows for various specifications to be examined. In line with Cioffi-Revilla [2010], a series of models are specified starting from the very simple. The simplest such model m^0, described in Equation 6, just adapts a schedule of age-specific fertility rates $\mu_f(x)$ according to a relative cohort-size measure C, in a similar way to is described in macro-models in the literature such as Wachter [1991] and Lee [1974]. This base-case model ignores the employment elements of the model, and allows a check that self-generated Easterlin-like cycles are indeed possible within the confines of the model as we have described it.

$$p(b_{xt}) = \mu_f(x)\exp(\eta C) \tag{6}$$

The model m^1 requires definition of aspiration levels for individual agents. These are defined with references to an agent's father's earnings at some formative age in their childhood (by default, at 15). Thus, in this case, an agent 'remembers' this value (denoted y), and an age-specific fertility function is adjusted according to a ratio of this value (averaged over the male and female members of a partnership) and the breadwinner's own earnings w_i. The log of this ratio replaces C in Equation 6 above.

For each of these models the age-specific fertility schedule is defined by a Hadwiger function [Chandola et al., 1999], as shown in Equation 7. The parameters of this function are fixed at values which allow a Total Fertility Rate (TFR) of just over 2 to be maintained, in absence of relative-cohort size adjustment. This means the population size stays relatively stable, making for easier analysis.

$$\mu_f(x) = a\frac{b}{c}\left(\frac{c}{x}\right)^{\frac{3}{2}}\exp\left(-b^2\left(\frac{c}{x}+\frac{x}{c}-2\right)\right) \tag{7}$$

These models provide a simplistic representation of the relationship between relative cohort size and income. However, they rely on adjusting an already calibrated fertility schedule according to relative cohort size. A better model would attempt to build the fertility from the bottom up.

An alternative model m^2 provides a different approach. Instead of allowing transitions between parities on a random basis, according to the ratio of income w to aspiration y, the model instead allows only a deterministic relationship, and a heterogeneous distribution of crucial decision-making parameters across agents provides for different behaviour across agents. Agents are endowed with an individual value γ_i determining what fraction below their parents' past earnings y_i they deem sufficient to start a family. Thus, the indicator $I_{i,1}(t)$ describing whether an agent has their first child in year t is defined as follows:

$$\begin{aligned} I_{i,1}(t) \text{ if } &: \\ & w_{it} > (1-\gamma_i)y_i \\ \gamma \sim \text{Uniform}(0, \gamma_{max}) & \end{aligned} \tag{8}$$

The individual preferences relating to consumption are uniformly distributed across the population, with a minimum value at 0 and the maximum determined by a variable parameter. Transitions to higher parities follow a similar process, except that we include both an adjustment for the notional cost of additional children as a proportion of income χ, and also allow for agents to limit their fertility once they reach their desired family size $par_i^{desired}$. The distribution of these preferences for numbers of children are determined according to another parameter. Equation 9 describes these conditions as the product of two indicator functions - on relating to relative income, the second to parity.

$$\begin{aligned} & I_{i,2+} \text{ if } : \\ & \qquad I\left\{(1-\chi(par)w_i) > (1+\gamma_i)y\right\} \; * \; I(par < par_i^{desired}) \end{aligned} \tag{9}$$

Once it has been determined that a couple decide to add to their family, a new agent is created. The newborn agent's characteristics are mostly initialised in their defaults states; age, experience, and (for females) parity, are for instance set to zero. Given the centrality of the concept of relative income, and the identification of one's parents as a comparator group, it is expected that social mobility and the correlations of income within families might impact on the fertility process modelled. The model therefore allows for intergenerational transmission of earning potential (for instance, through investment in education and parental attention) by allowing an agent's 'skill' value to be correlated with the average of it's parents $\bar{s}$. The degree of correlation is a parameter k in the model, so that:

$$\begin{aligned} \eta \sim N(k * \phi(\bar{s}), 1-k^2) \\ s_{child} = \Phi(\eta) \end{aligned} \tag{10}$$

where ϕ and Φ represent the standard normal density and distribution function respectively.

Initialisation

The model initialisation is a complex process. The starting population of 5000 agents receive random ages drawn according to a specified distribution, and random experiences commensurate with their age. The age distribution is chosen so that there are fewer agents at older ages, as might be expected because of mortality, and furthermore, the cohort aged between 40-60 at the start of the population is out-sized relative to those younger than it. Thus, those entering childbearing at the start of the simulation should belong to a small cohort, providing some impetus to the Easterlin-like effect that is the target for the simulation. Partnership and employment status setup amongst the initial population involves randomly assigning individuals into states for each age group (increasing with age), with probabilities designed to result in realistic population proportions. Those assigned to the 'married' and 'employed' states are matched with appropriate partners and jobs respectively. Similarly, children in the starting population are randomly assigned parents who have met their aspirations, which are again drawn at random. Some artefacts are evident in this setup, due to the difficulty of matching starting aspirations to as-yet unrealised partnerships and earning potentials, so a burn-in phase of 50 time-steps is allowed before simulation results are collected.

Model execution

The simulation class controls and coordinates the simulation. Time-step lengths are configurable, but one year steps are used in the result presented. Each time-step, agents are 'aged-on', and mortality, partnership search, job applications and fertility behaviour all take place, in that order. Following this, the matching in the marriage market is resolved, and similarly, job applications are assessed. The simulation code is designed so that the parameters for individual runs in an experiment can be distributed to independent instances of the program. This allows us to take advantage of the 'embarrassingly parallel' nature of the task of running multiple simulations (in that the results of each simulation are independent of each other), and utilise supercomputer resources such as the University of Southampton's `iridis 4` to minimise runtime. The simulation is run with a starting population of 5000 agents, but because some of the runs exhibit high levels of population growth and all run over 400 simulation-years, the final population size can be much higher.

Simulation Results

Simple probabilistic models

Starting with the simplest model m^0, simulation runs at suitably high values of the feedback parameter allow cycles in the times series of births to arise, as can be clearly seen in Figure 3. This is not surprising given the deterministic relationship between cohort-size and fertility in the model, and the work of Lee

[1974] and Wachter [1991], who show that such birth cycles are possible under at least some formulations of the theory.

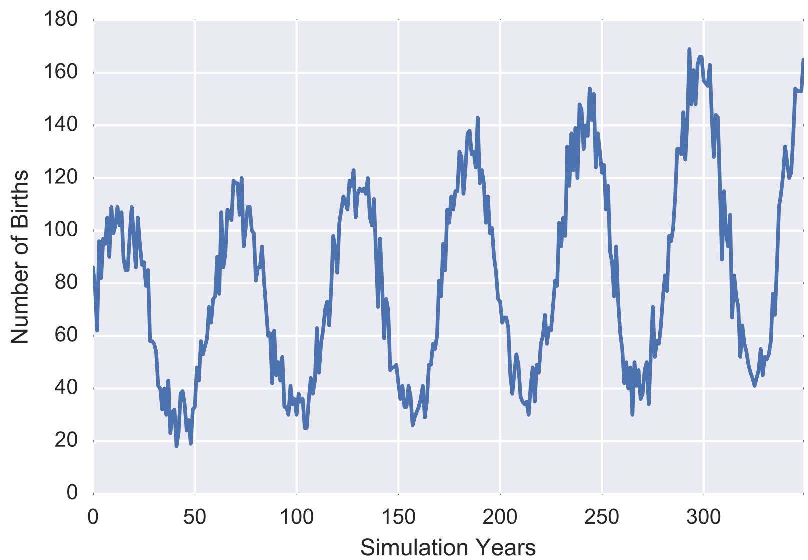

Figure 3: Plot of time-series of simulated births for model m0

By examining some other quantities it is also possible to see that although wages and unemployment are unrelated to the fertility process for this specification, they are clearly effected by changes in cohort size, which can be seen in Figure 4. This lays the groundings for later models, where fertility is affected by cohort size only through the labour market.

Moving on to the model m^1, which involves an adjustment of an age-specific fertility schedule based on relative cohort size, we note that the simulation is again able to recreate the expected cycles (Figure 5), giving some credence to the idea that relative earnings can form a plausible intermediary between cohort size and fertility.

However, for both these models, problems remain. The distribution over parity is particularly problematic, as can be seen in Figure 6. One would expect parity 2 or 3 to account for the highest share of families, but this is not the case. The lack of distinction between births of different orders in the decision making model explains the lack of realism in the parity distribution.

Furthermore, in this model, the age-specific fertility schedule is pre-fixed, and does not emerge from individual decisions. Thus, although the variation in this schedule results from individual level comparisons between own and parental incomes, a model in which the timing of childbearing is endogenous would be

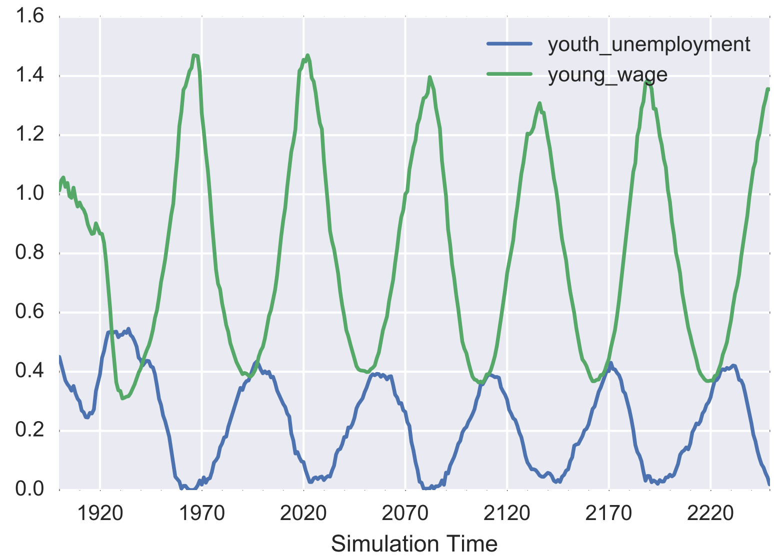

Figure 4: Plot of time-series of simulated youth unemployment and wages for model m0

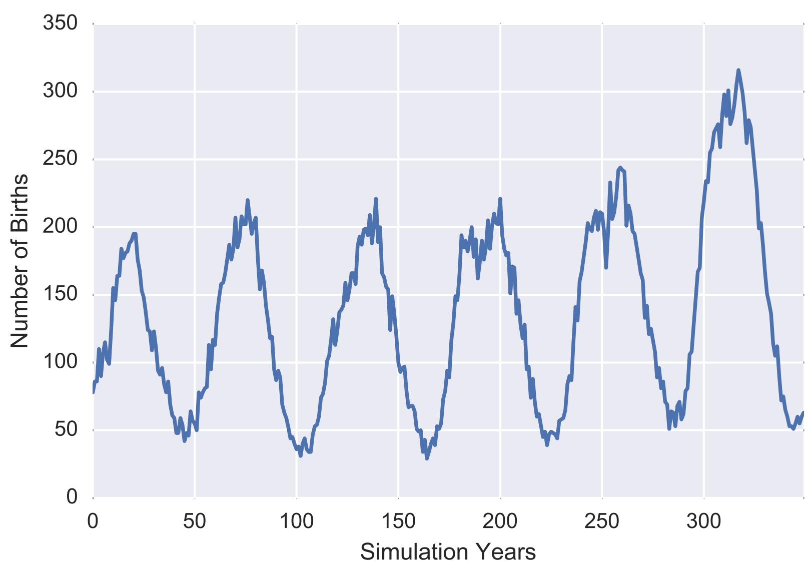

Figure 5: Plot of time-series of simulated births for model m1

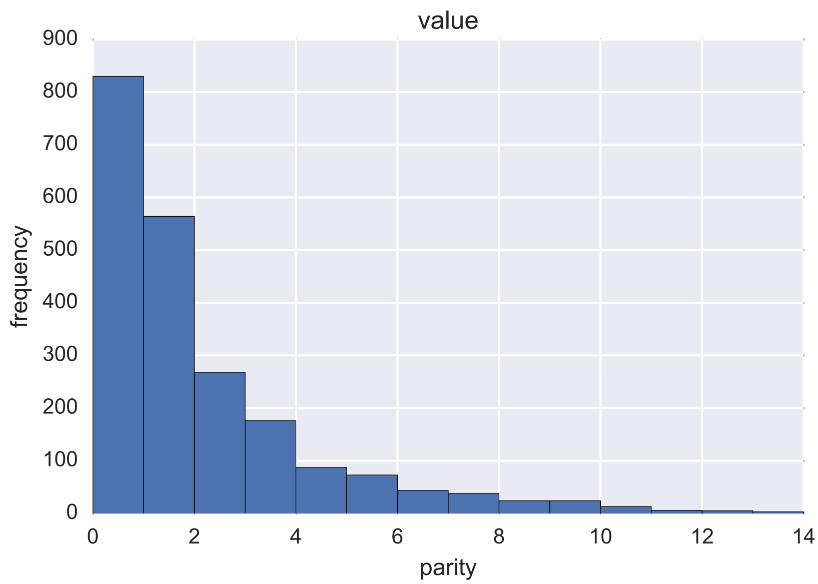

Figure 6: Plot of simulated parity distribution for a year in simulation m0

preferred.

Heterogeneous Agents Model

Model m^2 provides a simulation that is less dependent on stochastic response and more on explicit decision making. To understand the inter-relationships of the various parameters in the model and how they impact on fertility patterns, the use of statistical emulation techniques is advisable. The involve using semi-parametric techniques to fit a statistical meta-model (in this case based on Gaussian Processes) to the underlying simulation model [Kennedy and O'Hagan, 2001]. Additionally, the inclusion of a reasonable number of parameters in the model, and a greater degree of uncertainty about how those parameters will impact on fertility means that we must attempt to calibrate the model to reproduce the empirical phenomena we are interested in [ibid, Vernon et al., 2010]. In order to identify plausible parameter values that reproduce empirical patterns, a large number of simulation runs were undertaken at a spread of values. More specifically, a Latin Hypercube sample of 400 points, each of which were repeated 5 times to allow an estimation of the effect of simulation stochasticity generated from the Monte-Carlo trials in the simulation. Each run involved 400 simulated time-steps with a starting population of 5000 agents. Seven parameters were varied within these runs, as given in the list below. More parameters could have potentially been added to this list, such as those governing the productivity function in Equation 4, but these were left for future investigation.

1. *Support Ratio* : Defines the relationship between weighted population size and number of jobs.
2. *Small Family Desire* : The proportion of agents who desire 2 children or fewer. The remainder of agents are split between a desire for 3 or 4 children
3. *Wage Growth Rate* : This describes the year-on-year increase in wages, excluding the increase associated with agent's increased experience. These wage increases are assumed to be driven by increases in economic production, and represented by m in Equation 5.
4. *Wage Elasticity* : This describes the extend to which wages respond to changes to labour supply, operationalised by the relative size of an agent's birth cohort (Given by ζ in Equation 5).
5. *Intergenerational Correlation* : The extent to which individuals have similar abilities to earn as their parents, assumed to be driven by investment in education and parental time investment (described as k in Equation 10)
6. *Child Cost* : The cost as a proportion of income of raising an additional child (parameter χ in Equation 9)
7. *Relative income offset* : Defines the maximum of the distribution of individual parameters describing how an agent's fertility responds to relative wage (defined in Equation 8 as γ_{max}).

In this initial group of runs, the majority of simulations showed population decline, due to too few agents reaching their aspired standard of living, while some others runs saw explosive growth. Thus, a heteroskedastic emulator [Boukouvalas, 2010, Boukouvalas et al. [2014]] was fitted to the growth rates r defined by the standard growth equation $P_t = P_0 e^{rt}$, and an initial screening process identified ranges of the input parameters where the population is growing. A relatively ad-hoc method was used to achieve this goal; ranges of each parameter were selected for which moderate population growth was predicted, and a new Latin-Hypercube sample was defined within these confines. Cutting the range in this manner facilitates the subsequent calibration process, in which the history matching techniques described by Vernon et al. [2010] are applied in order obtain a subset of the parameter space for which the cycles of the desired period and amplitude can be obtained, combined with a similar rate of population growth.

History Matching of Easterlin Model

History Matching setup

To analyse subsequent waves of simulation runs, the time series of births from each simulation run in the repeated Latin Hyper-cube designs were first detrended by dividing by a trend estimated using LOESS smoothing, which fits a local regression function at each observation to a subset of points, with the points weighted by distance from that particular observation [Cleveland et al., 1998]. In common with Wachter [1991], we therefore observe the proportional deviation in births from the trend. Next, a non-linear model is fitted to the detrended series with the following form:

$$b*_t = 1 + \alpha sin(2\pi t\phi + \psi)$$

where $b*_t$ denotes the de-trended birth sequences, ϕ and ψ describe the frequency and phase of the time series respectively, while α is related to the amplitude. To find suitable starting points for the fitting process, a discrete Fourier transform of the detrended series is taken, and the dominant frequency f (equivalently, the dominant period $p = 1/f$) extracted from the spectrum. This measure of periodicity, together with the average rate of growth r, and the amplitude of the cycles are the important results from this process. According to Wachter [1991], a period of 42 years and a growth rate of 0.008, together with an amplitude of around 20% of the overall level of births, were seen in empirical data in the United States [ˆempirical]. We therefore wish to recreate these metrics in our simulated time series.

To this end, heteroskedastic emulators are fitted to three outputs; the log of the estimated period, the log of the amplitude, and the growth rate, and three iterative waves of history matching were carried out in order to restrict the parameter space to a set of points at which the targeted metric values were observed. The results from these points are described in the next section.

Results of Calibrated Simulation

From the remaining space of non-implausible parameters, 10 points are selected in order to examine the behaviour of the calibrated simulation for summary purposes. The points were chosen to maximise coverage of the remaining viable 2,700 parameter combination by generating a number of potential 10-point samples, and choosing the one which had the largest minimum distance between points. The simulation was run 7 times at the selected points, and a range of statistics were collected for each run.

To briefly characterise the calibrated points, the growth rates clustered around the targeted 0.008 rate, and the detrended series mostly had the desired periodic characteristics, as is evident in Figure 7. Some variability between the different time series is evident; some runs appear to display very strong periodic tendencies which remain constant throughout the simulation period. Others in contrast appear to be deteriorating and may be only transitory.

Other quantities aside from those calibrated against are also of interest, in particular in allowing some degree of validation of the simulated results. Figure 8 displays age-specific fertility for each of the 10 calibrated points over a period of 25 simulated years. This period is chosen in an attempt to capture a peak and a trough in fertility fluctuations. Also plotted (in red) are empirical ASFRs for the US for the peak year of 1957 and the earlier low point of 1936. As can clearly be

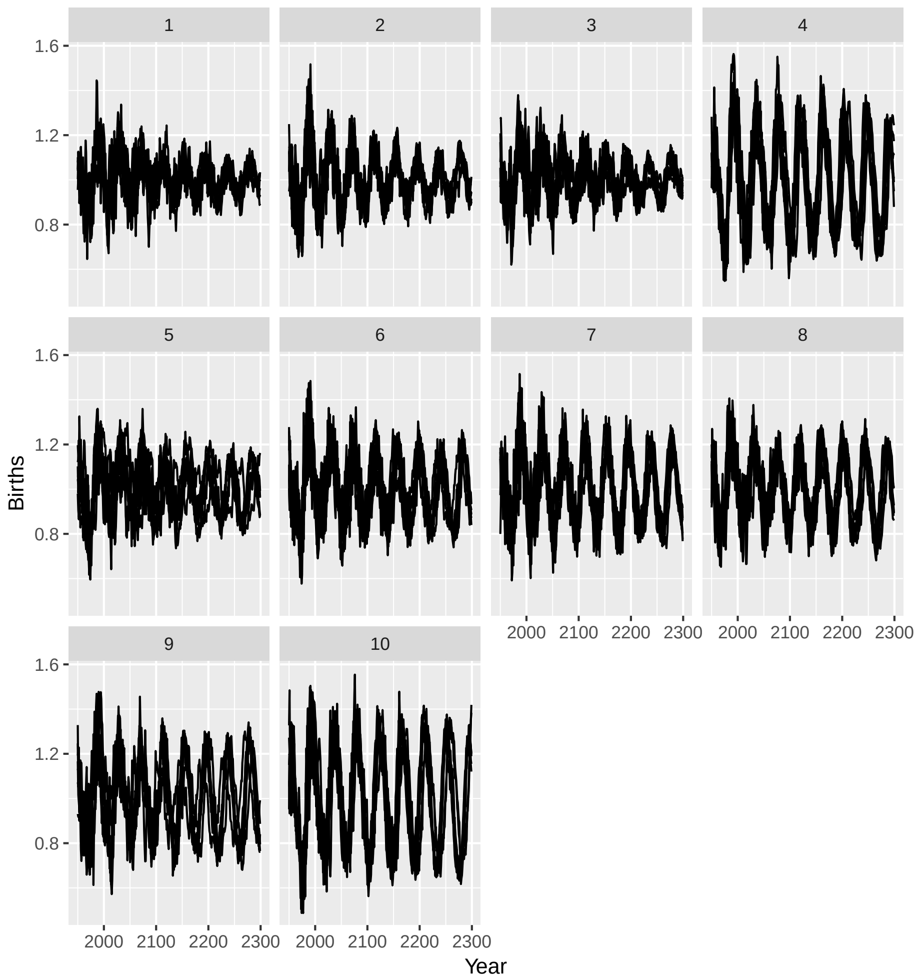

Figure 7: Simulated time-series of detrended births for a selection of non-implausible points

seen, the simulated fertility schedule has a peak in broadly the correct place, but the peak is more pronounced than in the empirical data, and declines too early. The shape of the fertility curve in the simulation is most likely related to the rate of growth of the productivity function with age. This helps define the rate at which individuals are able to catch-up with their parents earnings, and thus meet their aspirations. The parameters relating to this function were not included in the calibration exercise, and so future work might focus on calibrating against the rates explicitly and add the relevant parameters to the set to be varied.

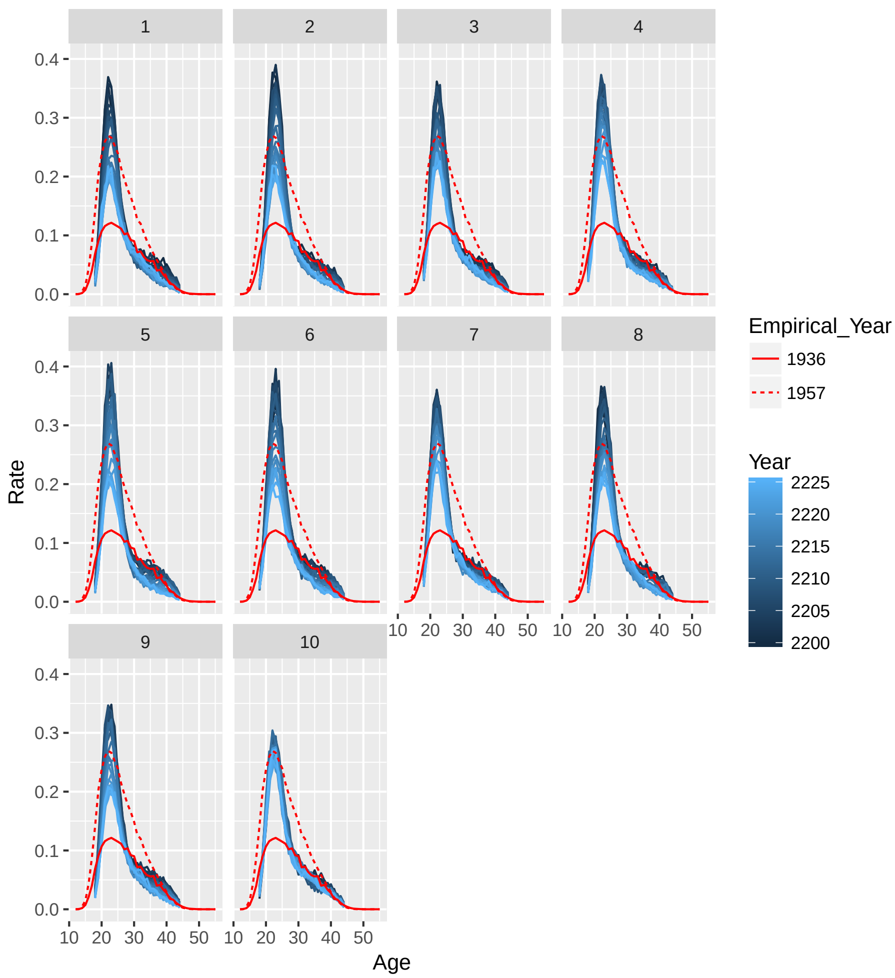

Figure 8: Comparison of Age-Specific Fertility Rates versus empirical US equivalents at peak and nadir of cycle. US Data from Human Fertility Database

Discussion

This paper identifies some conditions under which Easterlin-like cycles in fertility can be generated, under a number of assumptions needed to abstract reality into the simulation described above. It thus contributes an existence or plausibility proof, examining how the mechanisms described by Easterlin could indeed have led to the observed cycles. Simulation provides a useful tool with which to investigate this question as it allows us to replicate the process under discussion. In a sense, it provides a more detailed formalisation of the Easterlin Hypothesis than existing mathematical treatments by Lee and Wachter are able to provide, because specification of the 'control' process (whereby fertility is restricted amongst larger birth cohorts) is at the micro-level, rather than assumed and specified in terms of elasticities on fertility rates as a function of past births. From a methodological point of view, the value of using statistical emulators to examine and calibrate simulation models is underlined by the use of Gaussian process emulators to identify areas of the parameter space which recreate the two-generation cycles described by Easterlin. One benefit of the approach utilised here relative to optimisation-type calibration schemes is that a range of parameter combinations are identified, rather than a single 'best fit' point that ignores the fact that other parameter combinations may provide results well within a reasonable error, but might better reflect reality on other dimensions that are not observed or calibrated on [Oakley and Youngman, 2011].

Further work could take a number of interesting directions, particularly given the extensibility of the simulation framework. Firstly, various scenarios that have been posited to explain the demise of cyclic patterns in fertility from the mid-80s onwards may be tested within the simulated environment. For instance, female labour force participation maybe gradually increased, and likewise immigration may be introduced. It might be expected that this would result in a weakening in the relationship between birth cohort size and labour market success.

An interesting consideration is the relationship between population heterogeneity and randomness in simulations such as these. Part of the difference between models m^1 and m^2 is the reliance of the former on random trials, while the latter uses deterministic decision processes. The use of probability in model m^1 is here 'standing in' for heterogeneity in how much importance individuals place on their income relative to their wage. The model m^2 'explains' some of this randomness with a deterministic relationship and population heterogeneity, and is preferable in that it also takes into account the effect of dynamic sorting in a way that a simple probabilistic relationship would not - in a similar manner to the way that frailty models are better able to estimate hazards by accounting for differential population compositions over time.

Another direction might be to try and extend the simplistic model of decision making provided in the model to include a greater degree of time-awareness, more explicit calculation on the part of the agents as to the desirability of labour and family changes, and some explicit representation of the limits to the information available to agents [c.f. Gray et al.]. More sophisticated treatment of economic factors may also be possible. In particular, the introduction of savings at the individual level could open up interesting additional areas of research, as could the introduction of firms as explicit agents in the labour market.

At present, the simulation does not display complex emergent behaviour, which is not necessary to capture the nature of the target system. However, social interaction effects may also be included in future work. At present agents interact only through the labour and marriage markets. However, moving from a measure of relative income based on parents to one using social network peers as a reference group may be one way to examine this effect. In this framework, individuals are discouraged from starting a family if they feel the earn less than their friends. Alternatively, information about the benefits and costs of family formation may be passed from early adopters to peers, so that knowing parents may in fact make you less likely to want to become one yourself. Effects of this nature also allow the currently exogenous family size preferences to become endogenous norms, as in Aparicio Diaz et al. [2011], mutated by knowledge of peer's experiences.

However, we must beware of the 'kitchen sink' approach to simulation modelling, whereby the modeller attempts to include every possible facet and nuance of the target system [Axelrod, 2003]. Adding more and more detail will generally cloud our understanding of the simulation, as well as adding more and more potential sources of error. Keeping the simulation modular in design allows us to consider these extensions in isolation, allowing us to test the implications of each extension individually [Gray et al.].

References

Belinda Aparicio Diaz, Thomas Fent, Alexia Prskawetz, and Laura Bernardi. Transition to parenthood: the role of social interaction and endogenous networks. *Demography*, 48(2):559–79, may 2011. ISSN 0070-3370. doi: 10.1007/s13524-011-0023-6. URL http://www.ncbi.nlm.nih.gov/pubmed/21523518.

Robert Axelrod. Advancing the art of simulation in the social sciences. *Journal of the Japanese and International Economies*, 12(3):16–22, 2003. ISSN 10762787. doi: 10.1002/(SICI)1099-0526(199711/12)3:2<16::AID-CPLX4>3.0.CO;2-K. URL http://doi.wiley.com/10.1002/(SICI)1099-0526(199711/12)3:2{\T1\textless}16::AID-CPLX4{\T1\textgreater}3.0.CO;2-K$\delimiter"026E30F$nhttp://link.springer.com/chapter/10.1007/978-3-662-03366-1{_}2.

A Boukouvalas. *Emulation of random output simulators*. Phd thesis, Aston University, 2010. URL http://eprints.aston.ac.uk/15776/.

Alexis Boukouvalas, Pete Sykes, Dan Cornford, and Hugo Maruri-Aguilar. Bayesian precalibration of a large stochastic microsimulation model. *IEEE*

Transactions on Intelligent Transportation Systems, 15(3):1337–1347, 2014. ISSN 15249050. doi: 10.1109/TITS.2014.2304394.

Pierre Cahuc, Stephane Carcillo, and Andre Zylberberg. *Labor Economics*. MIT Press, Cambridge, MA, 2 edition, 2014.

T Chandola, D a Coleman, and R W Hiorns. Recent European fertility patterns: fitting curves to 'distorted' distributions. *Population studies*, 53(3):317–329, 1999. ISSN 0032-4728. doi: 10.1080/00324720308089.

C Cioffi-Revilla. A methodology for complex social simulations. *Journal of Artificial Societies and Social Simulation*, 13(1):7, 2010. URL http://jasss.soc.surrey.ac.uk/13/1/7.html.

W. S. Cleveland, E. Grosse, and W. M. Shyu. Local regression models. In J. M. Chambers and T.J. Hastie, editors, *Statistical models in S*, chapter 8. Wadsworth and Brooks/Cole, 1998.

Richard A Easterlin. The American Baby Boom in Historical Perspective. *NBER Occasional Paper*, (79), 1962.

Richard A. Easterlin. On the relation of economic factors to recent and projected fertility changes. *Demography*, 3(1):131–153, 1966. URL http://link.springer.com/article/10.2307/2060068.

Richard A. Easterlin. An economic framework for fertility analysis. *Studies in family planning*, 6(3):54–63, mar 1975. ISSN 0039-3665. URL http://www.ncbi.nlm.nih.gov/pubmed/1118873.

Richard A. Easterlin. What will 1984 be like? Socioeconomic implications of recent twists in age structure. *Demography*, 15(4):397–432, 1978. URL http://link.springer.com/article/10.2307/2061197.

Richard A Easterlin. *Birth and Fortune: The Impact of Numbers on Personal Welfare*. University of Chicago Press, Chicago, 2nd edition, 1987.

Richard A Easterlin. *The Reluctant Economist*. Cambridge University Press, Cambridge, 2004.

Joshua M Epstein and Robert Axtell. *Growing Artifical Societies; Social Science from the Bottom Up*. Brookings Institution Press; MIT Press, Washington DC; Cambridge MA, 1996.

G Fagiolo, G Dosi, and R Gabriele. in an Evolutionary Model of Labor Market and Output Dynamics. 7(2):157–186, 2004.

JC Frauenthal and KE Swick. Limit cycle oscillations of the human population. *Demography*, 20(3):285–298, 1983. URL http://link.springer.com/article/10.2307/2061243.

Jonathan Gray, Jason Hilton, and Jakub Bijak. Choosing the Choice: Refections on Modelling Decisions and Behaviour in Demographic Agent-Based Models.

Marc Kennedy and A O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society*, 63(3):425–464, 2001. URL http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00294/abstract.

Ronald Lee. The formal dynamics of controlled populations and the echo, the boom and the bust. *Demography*, 11(4):563–585, 1974. URL http://link.springer.com/article/10.2307/2060471.

Ronald Lee, Sang-hyop Lee, and Andrew Mason. Charting the Economic Life Cycle. *Population (English Edition)*, 34(2008):208–237, 2011.

Diane J. Macunovich. Fertility and the Easterlin hypothesis: An assessment of the literature. *Journal of Population Economics*, 11(1):53–111, mar 1998. ISSN 0933-1433. doi: 10.1007/s001480050058. URL http://www.springerlink.com/content/5n70pgeqbqwccmbg/.

Jacob Mincer. Schooling, Experience and Earnings. 1974.

Jeremy E. Oakley and Benjamin Youngman. Calibration of a Natural History Model of Bowel Cancer Progression Using Likelihood Emulation. *arXiv*, pages 1–24, 2011.

F C Pampel and H E Peters. The Easterlin effect. *Annual review of sociology*, 21: 163–94, jan 1995. ISSN 0360-0572. doi: 10.1146/annurev.so.21.080195.001115. URL http://www.ncbi.nlm.nih.gov/pubmed/12291060.

Nick Parr. Satisfaction with life as an antecedent of fertility. *Demographic Research*, 22:635–662, apr 2010. ISSN 1435-9871. doi: 10.4054/DemRes.2010.22.21. URL http://www.demographic-research.org/volumes/vol22/21/.

Michael S. Teitelbaum and Jay Winter. *The Global Spread of Fertility Decline: Population, Fear, and Uncertainty*. Yale University Press, New Haven, 2013.

Ian Vernon, Michael Goldstein, and Richard G. Bower. Galaxy formation: A Bayesian uncertainty analysis. *Bayesian Analysis*, 5(4):619–670, 2010. ISSN 19360975. doi: 10.1214/10-BA524.

Kenneth W. Wachter. Elusive Cycles: Are there Dynamically Possible Lee-Easterlin Models for U.S. Births? *Population Studies*, 45(1):109–135, mar 1991. ISSN 0032-4728. doi: 10.1080/0032472031000145116. URL http://www.tandfonline.com/doi/abs/10.1080/0032472031000145116.

Kenneth W. Wachter and Ronald D. Lee. U.S. Births and Limit Cycle Models. *Demography*, 26(1):99–115, 1989.

Brigitte Waldorf and Pillsung Byun. Meta-analysis of the impact of age structure on fertility. *Journal of Population Economics*, 18(1):15–40, 2005. ISSN 0933-1433. doi: 10.1007/s00148-004-0199-9. URL http://link.springer.com/10.1007/s00148-004-0199-9.

Sabine Zinn and Jan Himmelspach. Mic-Core: A Tool for Microsimulation. In Eds. M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, editor, *Proceedings of the 2009 Winter Simulation Conference*, pages 992–1002. IEEE, 2009. URL http://ieeexplore.ieee.org.