A Case Study: Intergenerational Fertility Effects

An Agent-Based Model of Intergenerational Fertility Effects

This paper aims to examine the plausibility of an existing theory which relates cycles in macro-level fertility to a individual-level desire to delay childbearing until a level of well-being commensurate with that achieved by one's parents [Easterlin, 1987]. 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 reviews 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 phenomena 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 allow a systematic analysis and calibration of the simulation described in the subsequent sections.

Background

The Easterlin Hypothesis attempts 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 [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 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'. Easterlin posits that generational cycles such as those described are the result of 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 likelihood that those in such cohorts will not feel well-off enough to start a family.

To breakdown this thesis further, I suggest it rests on four specific assumptions [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 latter statement contains a subsidiary assumption about the labour market, namely, that older and younger labour 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-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 temporal. 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 chapter focuses only on intra-cohort competition as it manifests itself through the labour market.

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 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].

Evidence for the hypothesis is mixed. At the macro-level, a considerable amount of work has been done creating formal models of the process, whereby fertility is effected 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).

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

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, evidence suggests that the existence of self-generating Easterlin cycles are possible, but that it is quite unlikely that the macro-forms specified are responsible for the patterns observed in the US [Wachter, 1991].

In terms of empirical data, it is clear that since the early 1980s the observed cycles have vanished, as can be seen in fig. 1. A number of reasons for this breakdown have been suggested. Increases in immigration and unemployment together with less secure working arrangements and more temporary work and increased female labour force participation have all been suggested [Pampel and Peters, 1995]. Indeed, the original analysis of Easterlin [1962] was very much concerned with the interplay of cohort size, immigration, and total labour force conditions in determining fertility, due to his initial association of fertility cycles with immigration-linked Kuznets waves in economic activity.

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, Macunovich [1998] believes that these negative finding are due in part to data problems and specification errors in the models used to test the theory. More recently, research has been conducted into the possiblity that happiness and satisifaction is often a prequisite 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 satisifaction.

A study by Waldorf and Byun [2005] attempted a similar review of macro-level studies in a more systematic manner using the tools of meta-analysis. Again, 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 substantially changed the direction of results, and highlighted endogeneity concerns in the use of this variable.

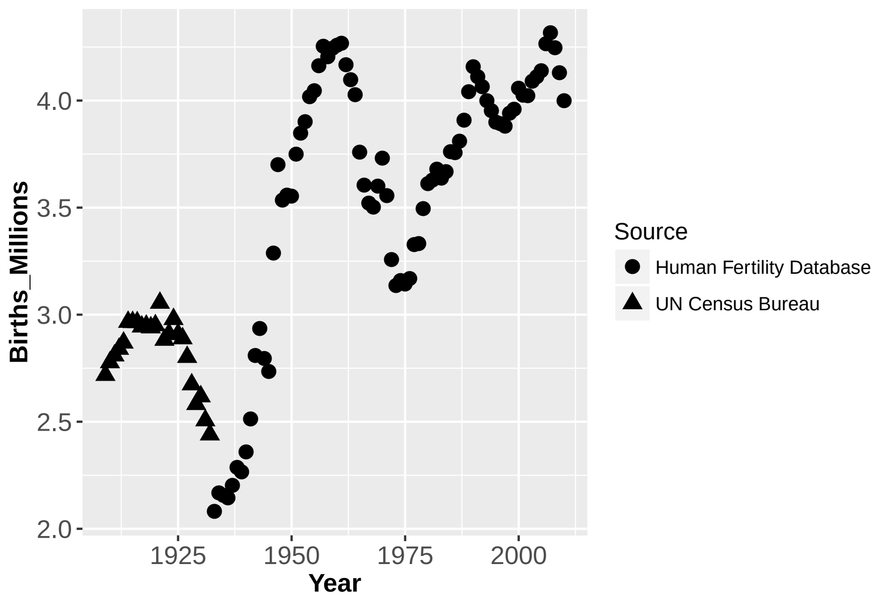

Figure 1: Time Series of US Births since 1909. Source: US Census Bureau, 1975; Human Fertility Database, 2015

Overall, there is no clear picture the literature as to whether Easterlin cycles where 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 forth [Macunovich, 1998]. This paper approaches the problem from the opposite direction, as will be discussed presently.

Research Questions

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 questions are given below

1. What behavioural rules and micro-level conditions are sufficient to generate Easterlin-like waves in fertility?

2. To what extent are such conditions realistic, and what does this tell us about the phenomena of 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 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 generation that results in the cycles observed in the data. This 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.

The logic of the 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 take involve parameters for which we do not know the true empirical value.

Model Specification

The model is coded in the programming language `python`, and consists of programmatic representations of two key process types; those relating to demographic changes and those relating to the labour market. Each agent possesses:

- An age
- A skill level drawn from a standard normal distribution
- A level of labour market experience, built from their employment history.
- A set of family relations, in particular, parents, and potential children and as spouse as well.
- A level of aspiration, set in childhood, determining at what level of well-being family formation is deemed desirable.

- A job (potentially). At present, the model is run with only male bread-winners in an attempt to reflect the reality of the period, although female labour market participation is possible within the simulation, and increases will be the modelled in future work. Jobs also have a difficulty level, which effects the wage to be gained in line with model in Cahuc et al. [2014], section 10.2.2.

Initialisation

The model initialisation is a complex process. The starting population has uniformly randomly drawn ages, and random experiences commensurate with their age. Partnership and employment status amongst the initial population drawing realistic proportions for each age group, and matching them with appropriate partners and jobs respectively. Similarly, children in the starting population are randomly assigned parents who have met their aspirations. Some artefacts are evident in this setup, so a burn-in phase of 50 timesteps is allowed before simulation results are collected.

In a typical time-step, an agent is subject to the following processes listed below. Timestep length are configurable, but 1 year steps are used in the result presented.

Ageing and mortality

Mortality is modelled simply using a Gompertz mortality function, with parameters set to reasonable values. These could be set empirically, but given that these are not key to the model, it was deemed that place-holder values were sufficient.

Partnership formation

Adult individuals are subject to a age-dependent hazard of partnership formation, again chosen to represent sensible transition times to partnerships. If agents are deemed by the simulation to experience the event, they are matched with those of the opposite sex in a similar position through a probabilistic matching function that pairs those close in skill and (offset) age variables, with provisions for a male-female age gap made.

Job Seeking

Once old enough, agents begin to apply for jobs. 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]. Agents apply at random to vacant posts, and offered jobs if they are the best candidate, assessed according to their productivity contribution, a function of their skill, experience and the difficulty of the job. Experience is assumed to affect productivity in line with the Mincer model of lifetime earnings [Cahuc et al., 2014]. That is, productivity increases with experience at a decreasing rate, before declining for higher values. Wages are determined related to productivity, but are adjusted in order to account for differences in labour supply between cohorts. More specifically, the adjustment factor is proportional to the size of their birth cohort relative to others currently in the labour market. Thus, the effect of supply of labour on wages is treated as exogeneous in this model. Agents accept the job if it is the best offer they receive. 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.

Aspirations

Aspirations are set directly based on parental income at some parametrically specified point in childhood. These are then the basis of family formation decisions.

Fertility

If in a partnership, female agents without children assess their well-being (household income) against their aspirations; if it is larger, they start a family. Subsequent fertility also depends on aspiration levels continuing to be met, but also on the duration since last birth. Additionally, fertility decreasing as parity increases (reflecting differences in prefered family size), and furthermore, to reflect changes in fecundity with age, the likelihood of giving birth declines at an increasing rate with age.

Simulation Results

In order to identify plausible parameter values that reproduce empirical patterns, a large number of simulation runs where undertaken at a spread of values. More specifically, a Latin Hyper cube sample of 400 points, several of which were repeated to allow an estimation of the effect of simulation stochasticity generated from the Monte-Carlo trials in the simulation. Each run involved 350 simulated timesteps with a starting population of 10000 agents. In order to assess whether Easterlin cycles do indeed exist, the time series of births from each simulation run was first de-trended using a LOESS smoothing algorithm. Next, a non-linear model is fited to the model with the following form:

$$b*_t = \beta_0 + \beta_1 exp(\lambda t) sin(2\pi t\phi + \psi)$$

where $b*_t$ denotes the detrended birth sequences, ϕ and ψ describe the frequency and phase of the time series respectively, β_1 is related to the amplitude, and λ captures the extent to which the waves in births are magnified or dampen with time.

To find suitable starting points for the fitting process, a discrete fourier transform is conducted, 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 are the important results from this process. According to Wachter [1991], a period of 42 years and a growth rate of 0.008 was seen in empirical data in the United States. We therefore wish to recreate these metrics in our simulated time series.

To investigate the parameter space, a Gaussian process emulator is fitted to the 'training data', in order to summarise the information gained from the 'training runs' above, accounting for model uncertainty [Kennedy and O'Hagan, 2001, O'Hagan, 2006]. This allows the prediction of simulation results at new points, as well as an assessment of the sensitivity of the model at a variety of parameters. The sensitivity analysis indicated that the most important parameters in the model were those relating to economic growth and the strength of wage response to supply of labour. Too much growth means that the parents' past earnings are not significant in determining fertility choices.

Results from one simulation run that adequately replicates the cycles in births in the US are displayed in Figure 2.

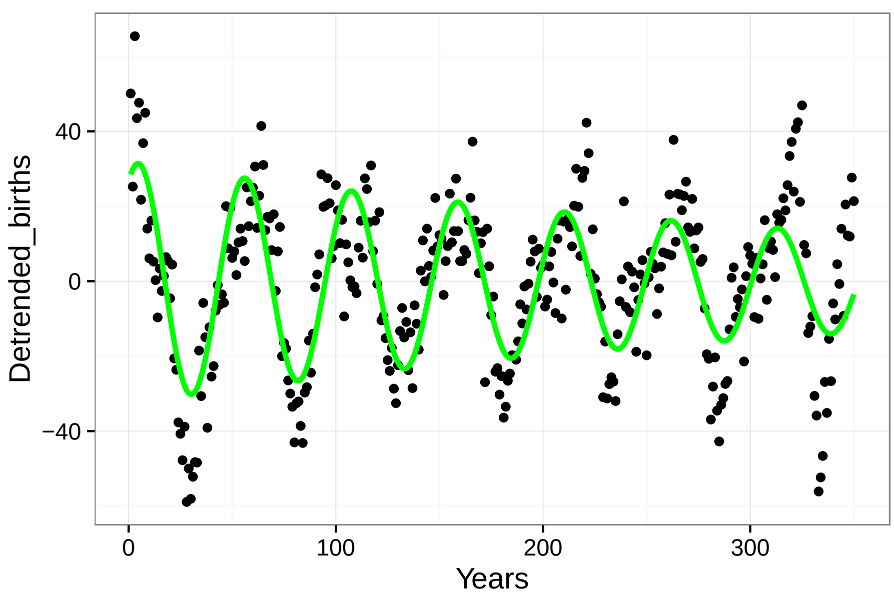

Figure 2: Detrended births with model fit

Discussion

This paper demonstrates 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 partly specified at the micro level, rather than assumed and specified in terms of elasticities on fertility rates as a function of past births.

Further work could take a number of interesting directions. 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 [Pampel and Peters, 1995]. It might be expected that this would result in a weakening in the relationship between birth cohort size and labour market success.

Another direction might be to try and flesh out the simplistic model of decision making provided in the model to include a greater degree of time preference, 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.

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.

Social interaction effects might also be included. At present agents interact only through the labour and marriage markets and through familial relationships. However, moving from a measure of relative income based on your parents to one using social network peers as a reference group may be one way to examine this effect [Billari et al., 2007]. 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.

However, we must beware of the 'kitchen sink' approach to simulation modelling. 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 only build in these extensions in isolation, allowing us to test the implications of each extension individually [Gray et al.]

References

Francesco Billari, Thomas Fent, Alexia Prskawetz, and Belinda Aparicio Diaz. The "Wedding-Ring": An Agent-Based Marriage Model Based on Social Interactions. *Demographic Research*, 17:59–82, aug 2007. ISSN 1435-9871. doi: 10.4054/DemRes.2007.17.3. URL http://www.demographic-research.org/volumes/vol17/3/.

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

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

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.

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: Some reflections on models of choice in agent-based modelling.

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.

R 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/.

Anthony O'Hagan. Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety*, 91(10-11):1290–1300, oct 2006. ISSN 09518320. doi: 10.1016/j.ress.2005.11.025. URL http://linkinghub.elsevier.com/retrieve/pii/S0951832005002383.

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/.

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.