DEMOGRAPHIC RESEARCH

A peer-reviewed, open-access journal of population sciences

DEMOGRAPHIC RESEARCH

VOLUME 29, ARTICLE 27, PAGES 729-766
PUBLISHED 9 OCTOBER 2013
http://www.demographic-research.org/Volumes/Vol29/27/
DOI: 10.4054/DemRes.2013.29.27

Research Article

Reforging the Wedding Ring: Exploring a Semi-Artificial Model of Population for the United Kingdom with Gaussian process emulators

Jakub Bijak

Jason Hilton

Eric Silverman

Viet Dung Cao

Table of Contents

Demographic Research: Volume 29, Article 27
Research Article

Reforging the Wedding Ring: Exploring a Semi-Artificial Model of Population for the United Kingdom with Gaussian process emulators

Jakub Bijak[1]

Jason Hilton[2]

Eric Silverman[3]

Viet Dung Cao[4]

Abstract

BACKGROUND
We extend the 'Wedding Ring' agent-based model of marriage formation to include some empirical information on the natural population change for the United Kingdom together with behavioural explanations that drive the observed nuptiality trends.

OBJECTIVE
We propose a method to explore statistical properties of agent-based demographic models. By coupling rule-based explanations driving the agent-based model with observed data we wish to bring agent-based modelling and demographic analysis closer together.

METHODS
We present a Semi-Artificial Model of Population, which aims to bridge demographic micro-simulation and agent-based traditions. We then utilise a Gaussian process emulator – a statistical model of the base model – to analyse the impact of selected model parameters on two key model outputs: population size and share of married agents. A sensitivity analysis is attempted, aiming to assess the relative importance of different inputs.

---

[1] Social Sciences, University of Southampton, Southampton, SO17 1BJ, United Kingdom. Corresponding author. E-mail: j.bijak@soton.ac.uk. Tel. +44 23 8059 7486.
[2] Social Sciences & Electronics and Computer Sciences, University of Southampton, UK.
[3] Electronics and Computer Sciences, University of Southampton, UK.
[4] Electronics and Computer Sciences, University of Southampton, UK.

RESULTS
The resulting multi-state model of population dynamics has enhanced predictive capacity as compared to the original specification of the Wedding Ring, but there are some trade-offs between the outputs considered. The sensitivity analysis allows identification of the most important parameters in the modelled marriage formation process.

CONCLUSIONS
The proposed methods allow for generating coherent, multi-level agent-based scenarios aligned with some aspects of empirical demographic reality. Emulators permit a statistical analysis of their properties and help select plausible parameter values.

COMMENTS
Given non-linearities in agent-based models such as the Wedding Ring, and the presence of feedback loops, the uncertainty in the model may not be directly computable by using traditional statistical methods. The use of statistical emulators offers a way forward.

1. Introduction

The aim of this paper is to reproduce selected illustrative features of population dynamics by using an agent-based model of a synthetic, closed population. Next to rules driving the agents' behaviour, the model additionally includes some real-world information, estimated and predicted from empirical data for the United Kingdom. In this paper we propose a method to explore and analyse selected properties of such models and identify plausible ranges of parameter space, by using Gaussian process emulators – statistical models of the underlying agent-based model.

Our overarching research goal is to try to explain the emergence of macro-level demographic patterns as a result of reasonable micro level assumptions which are explored in the model. We argue that models of this type, when backed with demographic data, can generate coherent 'what-if' scenarios based on a set of plausible assumptions of the mechanisms underlying the behaviour of simulated individuals (agents). In this way our approach combines the advantages of pure agent-based modelling, such as description of the mechanisms involved in population processes (Billari et al. 2007; Aparicio Diaz et al. 2011), with the empirical relevance of demographic analysis, when doing so is feasible. It has to be noted that our models are illustrative rather than attempting to be fully realistic with respect to all aspects of the underlying demographics. Some assumptions we make are simplistic, for the sake of transparency of the argument. Still, along the lines of earlier suggestions by Silverman, Bijak, and Noble (2011), we aim to enhance the predictive capabilities of the scenarios obtained, as compared to pure agent-based approaches.

In general, agent-based models (ABMs) are a class of computational models designed to simulate the interactions of autonomous agents which may represent individuals or groups. The goal of such models is to assess the effects of these actions on the overall system, or in other words to replicate an incidence of complex macro-level phenomena by simulating the actions of simple, micro-level agents (for detailed discussions relevant to social science, see Epstein and Axtell 1996, Gilbert and Tierna 2000, and Silverman and Bryden 2007). As a consequence, these simulations will generally include simple behavioural rules for autonomous agents, with the goal of observing how these low-level behaviours interact to produce higher-level complexity.

The simulations themselves can take a number of forms, ranging from very abstract models of a singular behavioural rule (as in Schelling 1971), to simulations of heterogeneous agents embedded in a geographical space (as in Axtell et al. 2002), to highly complex models incorporating agents with sophisticated neural networks for learning and decision-making (as in Hutchins and Hazelhurst 1995). In general, the particular form agents take in a given simulation varies significantly, and depends on the nature of the behaviours under examination. Indeed, the defining aspect of an agent-based model is not the particular representation of the agents, but instead the fact that the behaviour of those autonomous agents is explicitly modelled to examine its effects on the overall system. As Bonabeau (2002: 7280) notes[5]:

> “A number of researchers think that the alternative to ABM is traditional differential equation modeling; this is wrong, as a set of differential equations, each describing the dynamics of one of the system’s constituent units, is an agent-based model.”

The current article is devoted to demonstrating some advantages of a statistical analysis of ABMs, additionally equipped with selected empirical data series, in the demographic context. The paper is structured into five sections. After this Introduction, in Section 2 we present a brief background on micro-simulation methods in demography, including existing examples of agent-based models, and identify their main methodological gaps. As a way of filling these gaps, in Section 3 we introduce a Semi-Artificial Model of Population (SAMP) based on a reimplementation of the ‘Wedding Ring’ model of Billari et al. (2007). The prefix ‘Semi’ indicates that the model is a cross-disciplinary hybrid, comprised of agent-based and demographic components alike, and thus is not fully ‘Artificial’ in the sense of Artificial Life, seen as

---

[5] We are very grateful to an anonymous Reviewer for drawing our attention to this interpretation.

"life made by humans rather than by nature" (Langton 1995: ix). The presentation of SAMP starts from describing the general architecture of the original Wedding Ring model (Billari et al. 2007), followed by a discussion of empirical and projected demographic inputs, as well as of emulator-based methods for analysing the uncertainty in complex computational models.

Results of the simulations based on SAMP are shown in Section 4. We start from replicating the Wedding Ring model as proposed by Billari et al. (2007), and supplementing it with selected time series of actual and projected demographic data. Subsequently we extend the analysis to a formal quantification of sensitivity of two chosen outputs (population size and percentage of ever-married agents) to changes in selected parameters of the Wedding Ring. This allows for identifying plausible sets of assumptions for creating various scenarios of population dynamics. Finally, Section 5 offers a brief discussion of the results, followed by main conclusions and suggestions for further work. This paper complements an earlier prototype (Silverman et al. 2013), by offering several extensions of the proposed approach into directions more relevant for demography. The code for the current (second) version of the model is available from the OpenABM archive, at (http://www.openabm.org/model/3549/version/2).

2. Background: Micro-simulations and agent-based demography revisited

Existing demographic methods for simulating populations at an individual rather than aggregate level cluster together under the heading of micro-simulation. For the most part this class of simulation is defined by a concern with prediction, and through its use of empirical transition rates or waiting times to determine how individuals move from state to state. For these reasons this approach can be contrasted with the agent-based methodology introduced above (Murphy 2003; Spielauer 2007). In this paper, prediction is understood in a wider sense, not only as forecasting of the future quantities of interest, but also as estimating some of the historical patterns for which directly observable information is not available.

However, it needs to be stressed that in the social simulation literature, prediction is rarely thought to be the sole goal of the modelling process. Epstein (2008) mentions sixteen goals of modelling in scientific enquiry, from explanation of the underlying processes, through illumination of their different features, such as the process dynamics and the core uncertainty, to offering guidance for data collection, and engagement with the users of the models and the general public. Specifically in the context of agent-based modelling and how it is usually applied, explanation and aiding intuition of the underlying processes come to the fore (*idem*; see also Billari and Prskawetz 2003).

Computational micro-simulations involving demographic components have a long history. Orcutt (1957) advocated their use to better understand and predict aggregate socio-economic variables, and over the intervening half-century micro-analytic techniques have been used frequently for forecasting and informing policy (e.g., Orcutt, Caldwell, and Wertheimer 1976; Smith 1987). More recently, van Imhoff and Post (1998) identified several reasons why micro-simulation might be preferred to macro-level projection techniques in demography:

1. Micro-simulations are able to deal more easily with populations where the state-space is large, or where continuous covariates such as income are involved;
2. Micro-simulations allow richer outputs (in the form of e.g., simulated life histories);
3. Micro-simulations can easily deal with links between individuals.

These advantages mean that micro-simulations are particularly appropriate for policy problems relating to changes to, for instance, tax and benefit systems, as these require detailed, disaggregated information about future populations cross-classified by attributes of the individuals involved (Gilbert and Troitsch 2005; Willekens 2005). Additionally, flexibility in being able to try out different policy scenarios and to see the implications of each in a simulation is appealing in this context (Orcutt, Caldwell, and Wertheimer 1976). The demographic components of such models generally take a sample of the total population in question from a census or other micro-data source, and simulate their future demographic behaviour. A review of existing policy-focused micro-simulations can be found in Spielauer (2007).

Current developments in micro-simulation projections are focused around the synergies to be made in linking micro- and macro-level population models. The MicMac project is an example of such an endeavour, which attempts to link a deterministic macro-level simulation with a micro-simulation of individual life histories (Willekens 2005; Zinn et al. 2009). In a similar mode to forecasting applications, micro-simulation is also used to work out the implications of certain sets of observed rates for past or present population structure or individual life histories. Bracher, Watkins, and Santow (2003), for instance, used micro-simulation to examine the likelihood of newly married men and women in Malawi being HIV-positive, given prevalent infection and marriage rates. Similarly, Thomson et al. (2012) studied the implications for fertility of different levels of divorce and remarriage in developed countries based on data from France.

The unifying feature of the above models is that they are all heavily data-driven; individual behaviour is usually modelled not by rules or processes reflecting the actual mechanisms involved, but by empirical transition probabilities (for a discussion, see e.g., Murphy 2003). Given this heavy leaning towards incorporating empirical information, the subjects of such approaches as the micro-simulation analysis have been labelled by Courgeau (2012) as 'statistical individuals'. Unfortunately, as simulations become more complex and the state spaces of these models become larger, they require ever-greater quantities of data in order to define such probabilities, and collecting such data is time consuming, expensive, and in some circumstances may be completely impractical (Silverman, Bijak, and Noble 2011).

As well as the above policy-focused examples, micro-simulations have also been employed heavily in understanding changes in kinship networks – the 'linked lives' of different individuals. Aggregate population models give no method of linking individuals to each other, and thus micro-simulation has an important advantage in this respect for the demography of kinship (e.g., Wachter 1987, Wolf 1994, Murphy 2004). Several demographers have used micro-simulation to forecast the relative frequency of particular types of familial relationships in the future (e.g., Reeves 1987, Smith 1987). Wachter (1995, 1997), for instance, examined the availability of biological and step-children for US pensioners to 2030, seeing that declines in the former were to some extent compensated for by the latter. This particular micro-simulation has a number of interesting features. Firstly, it relies on a careful reconstruction of rates of population back to 1900 to generate kinship information for the present before projection could even begin. Secondly, it incorporates uncertainty surrounding future mortality directly into the simulation by including different stochastic realisations of Lee-Carter predictions in each replication. Thirdly, it maintains a linked 'offshore' population of potential immigrants which feeds into the main (US) population (although no further immigration is modelled after 1990 because of extreme uncertainty about the levels/rates).

This example of kinship demography micro-simulation, while clearly sophisticated, retains the data-dependency of other micro-simulation models discussed above. A slightly different approach, closer to agent-based modelling in its philosophy, is exemplified by work on anthropological and historical demographic contexts where data are often limited or non-existent. In these areas micro-simulation is used to test certain hypotheses about social and demographic phenomena, as opposed to working out the implications of certain combinations of rates (Wachter 1987, Hammel, McDaniel, and Wachter 1979). In contrast to most of the projection work described above, these models sometimes utilise behavioural or hypothetical rules instead of empirical transition rates to determine the actions of individuals in the simulation. For instance, Hammel, McDaniel, and Wachter (1979) tested the hypothesis that incest taboos might limit the viability of small human populations through the constraints placed on fertility. In their work micro-simulation methods are used to examine the implications of different severity of incest taboos formalised as behavioural rules that determine who a simulated individual can form a partnership with. Similar work has been undertaken for the inter-generational structure of households in pre- and post-industrial Britain, and how this depended upon individual decision rules about household formation (Wachter 1987, cf. Ruggles 1987).

Ruggles (1987) has critiqued approaches of this nature and claimed that hypothetical rules that formalise theories often depend on arbitrary choices by the modellers, as social theory is not described to the extent needed to incorporate it in a simulation. This criticism is perhaps misplaced: being forced to consider the implications of a hypothesis through formalising it can only help us understand it better, and the plausibility of different assumptions can be tested (Epstein 2008). Ruggles has also claimed that behavioural rules are either too simple to truly represent human behaviour, in which case hypotheses about social behaviour will tend to be disproved, or they are too complex for anyone to understand how they affect model outputs in practice (*idem*). However, Ruggles seems to have ignored two important possibilities: that one might vary the rules systematically to understand how they affect model results, and that the behavioural rules themselves might be the object of theorising.[6]

Drawing from this overview, we can make a distinction between two forms of individual-based simulation modelling. Micro-simulation aims primarily to predict, and utilises empirical probabilities to describe how individuals take on states and attributes, whereas agent-based modelling is concerned with explaining phenomena, and gives its agents rules that determine their behaviour. The difference between the two approaches is subtle and can be controversial: for example Murphy (2003) argues for including behavioural components and feedback mechanisms into the assumptions of micro-simulation models. From the perspective of social simulation, Gilbert and Troitzsch (2005: 13) see the distinction between micro-simulations and agent-based models mainly in the following areas:

- How many levels can be analysed: micro-simulation traditionally has been dealing with two levels (individuals and populations), while for agent-based models there could be more.
- Whether the simulated individuals communicate and interact with themselves and their environment, as is the case in agent-based approaches, which can explicitly bring the feedback mechanisms into the model.

---

[6] In later work, Ruggles has also raised concerns about the use of micro-simulation for the purpose of tracking kinship availability. The failure of most micro-simulation models to include correlation in demographic characteristics between parents and children, particularly for fertility, tends to mean that extreme kinship patterns are under-represented in simulated populations, and the variance of counts of available kin is underestimated (Ruggles and De Quincey 1993). Ruggles's solution is to collect further data to allow estimation of this correlation, extending the already onerous data requirements of such models.

- How many individuals are simulated: unlike in micro-simulation models, agent-based models represent scaled-down versions of the societies under study and thus are typically concerned with much fewer simulated individuals.

However, from the demographic point of view these distinctions are not clear-cut. There exist micro-simulation models of human populations, a notable example being the SOCSIM model developed at the University of California at Berkeley (http://lab.demog.berkeley.edu/socsim), which can include feedback effects and behavioural assumptions[7]. Moreover, many micro-simulation models are looking at more than two levels of analysis, with family or household-level structures coming to the fore, as discussed above. Hence, some of the apparent disconnection between these two approaches may be more due to different terminology used in the two disciplines – demography and social simulation – rather than to actual differences between the models used.

The existing examples of applying agent-based models in population-related applications are scarce, yet varied. From the classical example of the residential segregation model of Schelling (1978), other applications include marriage formation (Billari and Prskawetz 2003; Todd, Billari, and Simão 2005; Billari 2006; Billari et al. 2006, 2007; Hills and Todd 2008), family-related decisions with respect to parenthood transitions (Aparicio Diaz et al. 2011), migration (Heiland 2003; Kniveton, Smith, and Wood, 2011; Willekens 2012) and other forms of residential mobility (Benenson, Omer, and Hatna 2003), as well as overall household dynamics (Geard et al. 2013).

In general, agent-based modelling serves to compensate for some of the shortcomings of dynamic micro-simulation in certain types of modelling situations. The empirically-driven, data-heavy techniques of micro-simulation work well when examining the intricacies of complex pension or taxation systems, for example – a dependence on substantial amounts of data makes sense here, subject to appropriate use of Occam’s Razor when engaging in model construction. However, these approaches fall short when one wants to examine more general aspects of population change. The lack of interaction between entities and the inability to represent related elements such as social networks or the effects of spatiality mean that micro-simulations are typically unable to capture the full complexity of the processes underlying the developments we see at the macro level. Agent-based models can provide a platform for the study of these interactions, which can allow us to represent ‘linked lives’ – a critical component when studying the impact of relationships between individuals during the life course (Dannefer 2003).

---

[7] We are grateful to an anonymous Reviewer for drawing our attention to this.

Similarly, Entwisle (2007) has noted the potential for harnessing the power of ABMs to understand the importance of locality and space in population models. On the other hand, the existing gaps in agent-based modelling include a lack of predictive power, rendering them as best suited for quasi-predictive applications through coherent scenario generation (Epstein 2008). In that respect the current paper attempts to narrow the gap between the behavioural assumptions of agent-based models, aimed mainly at explanations and guiding intuition about phenomena, and higher predictive power of demographic micro-simulations. To do so, we propose a Semi-Artificial Model of Population which aims to bring the two methodological approaches closer together. The model is introduced in the next section.

3. Semi-Artificial Model of Population

3.1 Model architecture

We propose that a Semi-Artificial Model of Population (SAMP) is a simple multi-level and multi-state model of population dynamics, combining the statistical and agent-based modelling approaches. The model follows the life courses of simulated individuals (agents), who are subject to empirical patterns of fertility and mortality. For illustration we use time-varying data on age-specific birth and death rates for the United Kingdom (UK) for the period 1951–2009, and their further predictions yielded by Lee-Carter type models, following the example of Wachter (1995, 1997) for micro-simulations. The agent-based component is related to the process of marriage, and thus also household formation. For this purpose we use an adapted version of the ‘Wedding Ring’ model of Billari et al. (2007). Since, as mentioned in the Introduction, SAMP is intended to be illustrative and exploratory, we have omitted other demographic processes, such as migration, for the sake of transparency.

The overall architecture of the model is presented in Figure 1, which is a schematic representation of sample ‘linked lives’ and possible life-course trajectories of several agents. In terms of multi-level structure, SAMP operates at three levels: individuals (agents), households, and the whole population, with a direct bottom-up aggregation between these levels. Various technical aspects of the model are discussed in more detail in Sections 3.2 and 3.3, whereas Section 3.4 describes a framework for analysing uncertainty in such a model, based on the concept of Gaussian process emulators.

Figure 1: A schematic representation of linked lives and possible transitions in SAMP

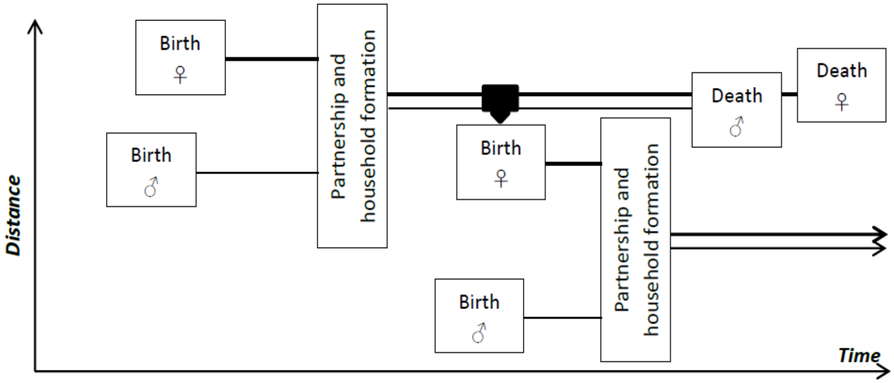

3.2 Agent-based component: Marriage formation on the Wedding Ring

In order to illustrate the potential benefits and pitfalls of combining the demographic micro-simulation and agent-based approaches, we have decided to replicate and expand upon the 'Wedding Ring' agent-based model of marriage formation designed by Billari et al. (2007). The model attempts to explain age-at-marriage patterns seen in contemporary developed countries. In summary, the Wedding Ring represents the process of marriage formation as a consequence of social pressure. Pressure arises from contact between married and non-married individuals within a given social network. This conceptual framework serves as a means of formalising some recent research in social influence and social learning, which has shown that these processes are highly relevant in individuals' decisions to get married (e.g., Bernardi 2003, *idem*). As reported by the authors of the Wedding Ring model, recent quantitative and qualitative research has attempted to investigate these effects further, and these studies appear to support the conclusion that social networks influence marriage and fertility decisions (Bernardi, Keim, and von der Lippe 2007; Bühler and Frątczak 2007; after Billari et al. 2007).

Thus the Wedding Ring model represents the spread of marriage through a population as a diffusion process, with social pressure from married agents influencing non-married agents within a given social network. However, as the authors mention, marriage differs from other diffusion processes, such as disease transmission, in that as well as experiencing diffusion pressure from ‘infected’ married individuals, marriage requires an individual to meet a suitable ‘uninfected’ (single) potential partner (Billari et al. 2007). The authors also note that recent marriages tend to have a higher impact on unmarried members of a social network than marriages in the distant past; this may be due to the tendency for the happiness increase provoked by a marriage to decrease over time, and thus presumably for the evangelist tendencies of recently-married couples to fade similarly as the years pass.

The Wedding Ring is so named due to the fact that in the original model agents live in a one-dimensional ring-shaped world (Billari et al. 2007). Each agent’s location is thus specified purely by a single coordinate (angle). The authors appear to have chosen the ring shape to avoid edge effects for agents located near a boundary. As the simulation progresses, each time-step in the simulated world is equivalent to one year. The agents are thus effectively situated in a cylindrical space, with one dimension of space and another of time (alternatively, age). Each agent’s network of ‘relevant others’ is then defined as a two-dimensional neighbourhood on that cylinder (*idem*).

Within that neighbourhood, the proportion of married agents determines the ‘social pressure’ felt by an individual agent, which influences their decision to seek out a partner (prospective spouse). The overall level of social pressure and the agent’s ‘age influence’ parameter determine the range in which agents search for suitable partners. The age influence value is defined using a piecewise-linear function which varied with the age of the agent. This function reflected the original authors’ assumptions that the size of an individual’s social network varies with age, peaking in middle age and decreasing rapidly thereafter. This function was retained in our reimplementation. As social pressure increases, agents widen their search range, and thus have a greater chance of successfully finding a partner (*idem*). However, the search is mutual: if one unmarried agent finds another within its acceptable range, marriage may only occur if the suitable partner has the searching agent within its acceptable range as well. Once married, agents may bear children; these children are then placed into the ring-world at a random spot in their parents’ neighbourhood and begin life at age zero. In the original model birth rates are continuously altered in order to keep the population size constant. This last restriction has been relaxed in our reimplementation, as has the limiting of childbearing to within marriage.

Agents in the original Wedding Ring model have a number of properties, primarily related to basic demographic statistics and to the implementation of the partner search process. Each agent is assigned a unique ID number at birth and their sex, year of birth and spatial location are recorded: since agents do not move in the Wedding Ring, the location remains constant. Over the life course further data are recorded for each agent, including their marital status, ID number for the chosen partner, marriage duration, relevant others, potential partner agents, length of the spatial interval for partner search, and derived social pressure.

In order to define the network of relevant others in which the agent searches for a partner, each agent is first classified at random into one of five possible types, according to which age ranges of agents they are most influenced by (i.e., similarly by younger and older agents, either mostly or only by older agents, or either mostly or only by younger agents). Then another random value is chosen which determines the scale of this age range, which in turn enables identification of the agent's possible relevant others. The size of the spatial interval for the agent's network of relevant others is symmetric around their location, and varies according to the size of the initial population; in our reimplementation we have included a parameter for 'spatial distance', denoted as *d*, which in this way indirectly determines the search space.

3.3 Demographic components: Mortality and fertility

Modelling of mortality and fertility in the Wedding Ring model is obviously a simplification of the overall mechanism of population dynamics. Given the tendency towards population ageing in modern societies, and given that older agents tend to have smaller networks of relevant others, failing to represent mortality in a more realistic fashion than by assuming constant rates could produce results that do not reflect the complexities of current population trends. Similarly, as populations continue to age, we see a significant drop in crude birth rates in many contemporary societies. With this in mind we incorporate empirical data on birth and death rates in the United Kingdom into our reimplementation of the Wedding Ring model.

To ensure that the starting structures within the simulation are reasonable, initial populations have been generated randomly but with agent distributions by age, sex, and marital status corresponding to the breakdown observed in England and Wales in the 1951 census.[8] To the same end, fertility and mortality rates experienced by agents over the course of the simulation are based on empirical and projected data for the United Kingdom. For mortality, the first 59 years of the simulation are based on age-specific mortality rates for the UK for 1951–2009. The data are split by individual year and single years of age from birth to the open interval 110+, and are based on population exposure estimates and death counts taken from the Human Mortality Database (2011).

---

[8] Source: Table 26 of the census output: [*Population by*] *ages (quinary) by marital condition*, by courtesy of the Office for National Statistics (ONS), Titchfield (personal communication on 29/11/2011). The structure for England and Wales was taken here as an approximation for the whole of the United Kingdom – not an unreasonable assumption given that, according to the ONS, in 1951 the population of England and Wales amounted to 82% of the total UK population.

To obtain rates for future years, in the horizon of the next half a century (2010–2061) predictions were produced using the well-known method developed by Lee and Carter (1992). The Lee-Carter mortality model reduces the age-by-time matrix of the logarithm of past mortality rates $\ln(m_{x,t})$ according to a time variant parameter k_t describing the overall mortality gradient; a vector a_x describing the average level of mortality at each age; and a vector of coefficients b_x showing how each age-specific rate changes over time relative to the overall mortality index k_t. Formally, this can be expressed as:

$$\ln\left(m_{x,t}\right) = a_x + b_x k_t + \varepsilon_{x,t}, \tag{1}$$

where $\varepsilon_{x,t}$ are normally distributed age-and-time-specific errors. Details of estimation procedure are available in Lee and Carter (1992). As only one parameter in this equation (k_t) varies with time, this simple representation of mortality is particularly amenable to forecasting through simple time series methods. In this application a Lee-Carter model was fitted to the data and the term k_t was projected forwards to 2061 using a random walk with drift. Future mortality rates for input into the simulation were then generated using the projected values derived from (1). The results of this exercise indicate a continual but slowing increase in life expectancy over the prediction horizon.

The base fertility rates were obtained in a similar way to those for mortality. Age-specific rates from 1973–2009 for UK women of childbearing age were obtained from the Eurostat database (Eurostat 2011), while earlier data for the period 1951–1972 were taken from the Office of National Statistics data for England and Wales[9]. No data were available for cohorts born earlier than 1920, so the small number of missing fertility rates for these cohorts during the period 1951–1965 were assumed to be the same as the earliest known rate for the respective age group. A Lee-Carter model for logarithms of age-specific fertility rates, $\ln(f_{x,t})$, was again fitted to the data, but, in contrast to the mortality predictions, two bi-linear terms $b_x k_t$ were required to best capture the trends in fertility (cf. Lee and Tuljapurkar 1994; Booth, Maindonald, and Smith 2002; Keilman and Pham 2006). The final model takes the form:

$$\ln\left(f_{x,t}\right) = a_x + b_{1x} k_{1t} + b_{2x} k_{2t} + \xi_{x,t}, \tag{2}$$

assuming normal errors $\xi_{x,t}$. An ARIMA(1,1,1) model has been then selected for each time-variant parameter k_{*t} in the above equation using standard selection procedures, as

---

[9] Source: ONS (1998). As with population structures, although the geographic coverage of the two sets of fertility data differs, the differences between the rates for England and Wales and those for the UK during 1973–1998 are negligible; thus the former are deemed a good substitute for the UK rates for 1951–1972.

implemented in the R package `forecast` (Hyndman 2011), and used to project future values to 2061. The resultant predictions see total fertility increase initially before converging at a value just above replacement level, and also display a continuation in the empirical trend towards later childbearing in the UK.

In order to ensure that fertility rates remain close to empirical values we also utilise empirical and projected values for the proportion of births to married mothers by year t and age of mother x, denoted here as $r_{x,t}$. The rate of childbearing for a simulated married woman is then calculated by taking the product $r_{x,t} f_{x,t}$ and multiplying it by the ratio of the total female population to the number of married women in a particular age group, for a given year:

$$f_{x,t}^{M} = r_{x,t} f_{x,t} \times \frac{P_{x,t}}{P_{x,t}^{M}} \tag{3}$$

In equation (3), superscript M denotes the population of married agents, and $P_{x,t}$ refers to the total simulated female population at age x and time t. Similar calculations are made for unmarried women's fertility by using the value $(1 - r_{x,t})$ and the ratio of the total population to the number of unmarried women. These calculations allow for childbearing amongst both married and unmarried simulated women in accordance with empirical and forecasted patterns, while also ensuring that total fertility remains broadly in line with overall expectations.

In order to calculate fertility rates for married and unmarried agents, then, values for $r_{x,t}$ were required for the whole range of the simulation 1951–2061. The Eurostat database (2011) provides data for numbers of births to married and unmarried mothers from 1960 to 2011, but splits this data by age of mother only between 1982–2010, allowing the calculation of $r_{x,t}$ for these years only. In order to get estimates for the remaining years, several steps have been taken. Firstly, another Lee-Carter model (4) has been fitted to the logit-transformed values of $r_{x,t}$ between 1982–2010:

$$\text{logit}\big(r_{x,t}\big) = \alpha_x + \beta_x \kappa_t + \zeta_{x,t}, \tag{4}$$

again with normal errors $\zeta_{x,t}$. The time varying element of this model κ_t is considered to be approximately proportional to the values of r_t, the proportion of the births to married women irrespective of age. Eurostat data for r_t for 1960–1981 could therefore be transformed in order to continue the times series for κ_t by subtracting the mean value of r_t between 1982–2010 and multiplying by the ratio of the standard deviations σ_κ / σ_r. Secondly, with estimates of values for κ_t from 1960–2011, this time series has been predicted backwards and forwards to obtain estimates for the periods 1951–1959 and 2012–2061. The auto.arima method of the `forecast` package (Hyndman 2011) was used to select a ARIMA(3,2,0) model for the backward prediction, although for predicting forwards the model was constrained to a single differencing in order to keep the values of κ_t to within a plausible range at the end of the forecast horizon, resulting in the choice of an ARIMA(1,1,2) model. Finally, with the computed values of κ_t and the age specific means and growth rates α_x and β_x, the complete matrix $r_{x,t}$ has been estimated from the equation (4).

3.4 Framework for analysing uncertainty: From Monte Carlo to Gaussian process emulators

The outcome of a single run of an agent-based model denotes one possible scenario of the demographic development of the artificial population under study. Thanks to its construction, such a scenario remains consistent at the macro level with the micro-level assumptions on the behaviour of individual agents. In general, scenarios are potentially very useful for answering many policy-relevant ‘what-if’ questions, by analysing the model responses to a changed setting of assumptions (parameters). On the other hand this scenario is only one of many possible under different parameterisations of the model, and even though in SAMP it is characterised by an increased degree of plausibility due to including observed demographic parameters it does not provide enough information on its predictive accuracy. As with all models aimed at predictive applications, some assessment of uncertainty is thus needed.

Due to the inherent non-linearities of relationships within agent-based models such as SAMP, and the presence of various feedback loops, the uncertainty of model outputs may not be easily computable analytically. Instead, a Monte Carlo simulation can be performed, where the model based on a pre-defined set of parameters is run many times, and the empirical realisations analysed in the form of statistical distributions. This solution would be appropriate for assessing the code uncertainty, related to variation in the realisations of the model itself (cf. O’Hagan 2006). Examples of applying the Monte Carlo approach to SAMP are presented in Section 4.3.

However, the code uncertainty is not everything. Considerable uncertainty is also associated with the unknown parameters driving the model assumptions. In principle, this issue could be also addressed using a Monte Carlo approach, although, given the potentially high dimensionality of the problem, the number of required iterations, coupled with the computational complexity of the models, means that the time required to run them quickly becomes prohibitive (Kennedy and O’Hagan 2001). An alternative approach is to construct an emulator – effectively, a statistical model of the underlying complex computational model, reduced to the inputs and outputs of immediate interest – and to examine its properties (Oakley and O’Hagan 2002). In order for the uncertainty of the emulator to be described coherently and correctly, the preferred underlying statistical framework is the one of Bayesian inference (*idem*).

Amongst methods that have been proposed for building emulators, the one that is argued to be relatively simple, yet very flexible in application to complex computational models, is based on Gaussian processes. A succinct introduction to Gaussian process emulators is provided below. In general, the theoretical foundations have been laid out, for example, in the work of O'Hagan, Kennedy and Oakley (see Kennedy and O'Hagan 2001; Oakley and O'Hagan 2002; Kennedy 2004; O'Hagan 2006), and on the website of the research community *Managing Uncertainty in Complex Models* (http://www.mucm.ac.uk), where the construction and estimation of such emulators is presented in much more statistical detail.

Thus, let function $f(\cdot)$ denote the base computational model of interest – in our case, SAMP. For the purpose of building an emulator the focus is on a pre-defined vector of n inputs, $\mathbf{x} \in \mathbf{X} \subset \Re^n$, and a single output, $y \in \mathbf{Y} \subset \Re$, such that $y = f(\mathbf{x})$. It has to be noted that $\mathbf{X}$ does not have to exhaust the whole parameter space of the underlying model, but rather should relate to those inputs (parameters) that are considered important from the point of view of the output studied. Following Oakley and O'Hagan (2002: 771) and Kennedy (2004: 2), we define a Gaussian process emulator, conditionally on a set of its parameters, as a multivariate Normal distribution for p realisations of the function f, $y_1 = f(\mathbf{x_1})$, … , $y_p = f(\mathbf{x_p})$, denoted jointly as $\mathbf{f}$ (*idem*):

$$[\mathbf{f}(\cdot)|\boldsymbol{\theta}, \sigma, \mathbf{R}] \sim N[\mathbf{m}(\cdot), \sigma^2 \boldsymbol{c}(\cdot,\cdot)] \tag{5}$$

The mean of the process, $\mathbf{m}$, is modelled through a vector linear regression function of $\mathbf{x}$, $\mathbf{h}(\mathbf{x})$, with coefficients $\boldsymbol{\theta}$ such that, for every output $f(\mathbf{x})$, $m(\cdot) = \mathbf{h}(\cdot)^{\mathrm{T}} \boldsymbol{\theta}$. Further, σ^2 is the joint variance parameter, and $\boldsymbol{c}(\cdot,\cdot)$ denotes a correlation matrix, the elements of which are here assumed as $c_{ij}(\mathbf{x_i}, \mathbf{x_j}) = \exp\{-(\mathbf{x_i} - \mathbf{x_j})^{\mathrm{T}} \mathbf{R} (\mathbf{x_i} - \mathbf{x_j})\}$. The diagonal matrix $\mathbf{R} = \mathrm{diag}(r_1, \ldots, r_n)$ is composed of *roughness* parameters $\{r_1, \ldots, r_n\}$, which indicate how strongly the emulator responds to particular inputs (Kennedy and O'Hagan 2001: 432–433; O'Hagan 2006).

In order to estimate the parameters of the emulator, simulation data $\mathbf{D} = [f(\boldsymbol{\delta}_1), \ldots, f(\boldsymbol{\delta}_N)]$ are required for a set of N experimental points $\boldsymbol{\Delta} = \{\boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_N\}$, where $\boldsymbol{\Delta} \subset \mathbf{X}$ (Kennedy 2004: 2). Making additional assumptions on the prior distributions of the parameters of the emulator (5) allows for applying full Bayesian inferential mechanism to obtain the posterior distribution of $\mathbf{f}$ given $\mathbf{D}$. In order to incorporate the code uncertainty into the emulator, an additional variance term (referred to as a nugget) can be subsequently included in the estimation of the mean and the covariance matrix of the posterior distribution (*idem*).

The emulator, once built, can be used for a basic uncertainty analysis, which looks at how much uncertainty in the output is being induced by the set of inputs **X** under study, treated here as random variables with some assumed probability distributions (e.g., Kennedy 2006). Subsequent uses may involve a probabilistic sensitivity analysis, whereby the impact of particular inputs on the output is assessed based on the reductions of the output variance due to actually observing particular inputs (for a theoretical discussion, see Oakley and O'Hagan 2004). Here, output variance reductions obtained by conditioning on true observed values of single inputs are referred to as main effects, and the additional reductions obtained for combinations of inputs as joint (interaction) effects. As an illustration, in Section 4 we present an example of an uncertainty and sensitivity analysis performed with a Gaussian process emulator for selected input parameters in SAMP, focusing on the share of agents who have ever entered into marriage, and on the overall simulated population size.

4. Selected results

4.1 Model implementation

SAMP was implemented in Repast Simphony v. 2.0, a Java-based environment especially designed for agent-based modelling and simulations. Each run of the model included 110 time steps, which in our case correspond to calendar years, starting with 1,600 agents in the simulated year 1951. The starting period was chosen in order to match the initial population structure with the 1951 census for the UK, as discussed in Section 3.3. The results presented in this section focus on the simulated year 2011, for which empirical verification of some aspects of the simulation was possible, and on 2061, which is the simulation horizon. The latter is mainly included to check the plausibility of trajectories of longer-run population dynamics, discussed in Section 4.3.

In general, the simulation has not proven to be too expensive in terms of computing time, with about 2–3 minutes per run on a standard desktop computer. Within each step of each run the updating is done in the following sequence: all agents are aged one year, then adult agents outside marriage undergo the partner search routine, female agents in the relevant age groups are subject to a random risk of childbearing, and all agents are subject to a random risk of dying. In line with its general aim, SAMP is intended to show the effects of adding dynamic fertility and mortality change to the Wedding Ring model of Billari et al. (2007), rather than being aimed at closely reproducing the UK population.

Two of the sub-steps of the modelled process of population dynamics are driven by empirical birth and death rates, either observed for simulation years 1951–2009, or projected using the mean forecast from the Lee-Carter model for the rest of the simulation period (2010–2061). The empirical fertility rates are adjusted to reflect differences in fertility patterns of agents within and outside marriage, as discussed in Section 3.3.

In the final part of each step, agents who have died are removed, newborn children are added to the population, and the agents' records are updated. The summary statistics are produced every simulated year, and refer to population structures and marriage hazard functions. The outputs also form a basis for building statistical emulators discussed in Section 4.2, which can be subsequently used to aid the construction of sample 'what-if' scenarios, such as the one presented in Section 4.3. In our example the construction of emulators is based on Δ consisting of 7^3 = 343 model runs, corresponding to seven design values for each of the three parameters considered, as detailed in Section 4.2. In the scenario exercise, for the purpose of the Monte Carlo analysis the model has been run 500 times for one given parameter set, in order to assess the uncertainty of population size resulting from the inherent randomness embedded in SAMP.

When re-implementing the Wedding Ring model and switching to empirical and projected vital rates, the original parameter settings of Billari et al. (2007) were no longer producing results that could be considered fully plausible in the light of the empirical evidence, as discussed further in Sections 4.2 and 4.3. Most importantly, this concerned the two parameters, α and β, related to the social pressure function $s(r)$, defined in the original paper as (Billari et al. 2007: 66):

$$s(r) = \exp\{\beta(r-\alpha)\}/[1+\exp\{\beta(r-\alpha)\}], \tag{6}$$

where r denotes the proportion of agents with partners within one's network of relevant others. The parameters were originally benchmarked as $\alpha = 0.5$ and $\beta = 7$ (*idem*).

In order to identify the range of plausible values of α and β for the purpose of setting coherent scenarios of population dynamics, two pre-simulation emulators (which we will call *premulators*[10]) were built first, following the discussion offered in Section 3.4. The details of their construction, and the results of the emulation, are presented next.

4.2 Uncertainty and sensitivity analysis: Population size and marriage rates

In this section we present two prototype Gaussian process emulators for SAMP, with the aim of identifying areas of the parameter space that result in empirically plausible

---

[10] In this context, the term was coined by the second author of this paper.

population dynamics and marriage processes. The focus here is on two features of the mechanism of the marriage formation: social pressure and spatial distance, both of which feed into the intensity of the partner search. In the first emulator we analyse the impact of the three underlying parameters: α and β in equation (6), as well as the spatial distance parameter d, on the uncertainty in the resultant overall share of population aged over 16 years who have entered into marriages at the simulation year 2011, denoted as p. Since p is bounded between 0 and 1, in order to preserve these constraints we have logit-transformed the output variable into $u = \ln[p/(1–p)]$.

To obtain the simulation data **D** for building the emulator, we have run the model on a Cartesian product of pre-selected input values, $\mathbf{\Delta}=\boldsymbol{\alpha}'\times\boldsymbol{\beta}'\times\mathbf{d}'$, where $\boldsymbol{\alpha}'= [0, 0.333, 0.666, 1.0, 1.333, 1.666, 2.0]^T$, $\boldsymbol{\beta}' = [\exp(–1), \exp(0), \exp(1), \exp(2), \exp(3), \exp(4), \exp(5)]^T$, and $\mathbf{d}' = [5, 10, 15, 20, 25, 30, 35]^T$. Subsequently, a basic sensitivity analysis of the output u to the variation in the inputs has been attempted, with the aim to assess the importance of the three parameters affecting the modelled process. For the sake of the clarity of the presentation we focus here on the period with observed rather than projected birth and death rates, the addition of which would necessitate the inclusion of additional sources of uncertainty in the analysis.

The emulator was constructed, and the uncertainty and sensitivity analysis was performed in version 1.1 of the dedicated software GEM-SA (Gaussian Emulation Machine for Sensitivity Analysis), written by Marc Kennedy and Anthony O'Hagan (Kennedy 2004; O'Hagan 2006)[11]. Since for the Gaussian process emulators defined in Section 3.4 analytical derivations of the posterior distributions are available conditionally on the estimated values of **R** (Kennedy 2004), performing a numerical simulation such as Markov chain Monte Carlo (MCMC) was not needed. Hence, a single run of the algorithm took only about two minutes on a standard desktop computer. The quality of the emulator construction has been assessed by using a leave-one-out cross-validation method, whereby the errors for single design points from $\mathbf{\Delta}$ are calculated by comparing the output values with the ones yielded by emulators based on all the remaining design points. The root mean-squared standardised error (RMSSE) reported by GEM-SA in this case equals 1.35, which indicates a strong emulator fit, in comparison with the ideal outcome of 1 expected for a perfectly fitting model.

In GEM-SA, the distributions for the parameters of the Gaussian process (5) are *a priori* assumed to be vague, with $p(\boldsymbol{\theta},\sigma^2) \propto \sigma^{-2}$ denoting limited information about the features of the process prior to observing the simulation data (inputs and outputs). The independent prior distributions for particular elements of the roughness matrix, r_i, are in turn exponential, with parameter $\lambda = 0.01$ (Kennedy 2004: 2). It can be shown (*idem*: 3) that this posterior distribution has a marginal multivariate Student's t density, conditional on the roughness matrix **R**, here estimated as **R** = diag(15.346, 19.417,

---

[11] The software is available from http://ctcd.group.shef.ac.uk/gem.html (retrieved on 15/07/2012).

0.277). For the purpose of the uncertainty and sensitivity analysis, in this application the three input parameters are assumed to be unknown and described by the following Normal distributions: α ~ N(1.0, 0.25), β ~ N(2, 2.25) and d ~ N(20, 56.25). These distributions were derived by assuming that the input parameter range was assumed to cover +/– two standard deviations from the mean. We have additionally assumed that the mean of the Gaussian process (5) is a linear function of its inputs. The code uncertainty has been handled by adding an additional error term (nugget) in calculating the posterior estimate of the covariance matrix.

The outcomes of the uncertainty analysis performed by using the emulator indicate a mean percentage of ever-married agents of p = 6.2%, corresponding to the logit-transformed variable u = 0.585. The variance parameter σ^2 is estimated as 0.504, and the nugget variance as 0.077, indicating that, for u, the uncertainty in the three inputs considered is much more important than the code uncertainty resulting from the randomness in the model. The total output variance in u induced by input uncertainties is estimated as 2.566, of which the emulator-related component contributed 0.0013. In terms of the sensitivity analysis the most important variables in the analysed setting proved to be the two parameters of the social pressure function, α and β, respectively accounting for 35.0% and 44.5% of the variability of the output, and their interaction contributing further 18.1%. The spatial distance parameter d was responsible for explaining only 1.1% of the variability of u.

A second emulator has been constructed taking population size simulation year 2011 (N) as an output, with the same input values as for the emulator for the percentage of ever-married agents. To ensure the emulator outputs remained bound from below by 0, the data were log-transformed as $M = \ln(N)$. The uncertainty analysis based on this emulator estimates the mean value of M as 7.58, corresponding to a population size N of 1959 agents. This time the parameter σ^2 is estimated as 0.990, and the nugget variance as 0.688, suggesting that the code itself is almost as important as the uncertainty in the underlying marriage formation process. The total variance on M is estimated as 0.00089, with 0.00002 being accounted for by the emulator-related component. Cross-validation statistics indicate that the fit of this emulator is worse than for the previous one, with an estimated RMSSE of 5.21, reflecting much larger variation across the design space Δ, which is not surprising given the important role of code uncertainty. In turn, the estimated roughness parameters are $\mathbf{R}$ = diag(0.967, 1.480, 0.669), suggesting that the emulator smoothens the underlying space to a suitable degree. The sensitivity analysis reveals the proportions of the variance accounted for by α as equal to 19.1% and β to 28.7%, with a further 7.3% of the variance being accounted for by their interaction. The spatial distance parameter d is more important than for the previous emulator, accounting for 37.5% of total variance.

As an illustration, Figures 2 and 3 show contour maps of the predicted emulator means for the output variables p and N, respectively, plotted against two parameters of the social pressure function (6), $\alpha \in [0, 2]$ and $\beta \in [\exp(-1), \exp(5)]$, and conditional on the constant value of $d = 25$. Despite relatively large uncertainty of the emulator for the percentage of ever-married agents illustrated in Figure 2, its predicted mean (64.2%) seems plausible and very close to the empirically observed value of 65.2% ever-married UK residents aged 16 or above, recorded in the 2011 census data for England and Wales provided by the Office for National Statistics (ONS 2012a).

Figure 2: Mean share of ever-married agents by social pressure parameters, simulation year 2011

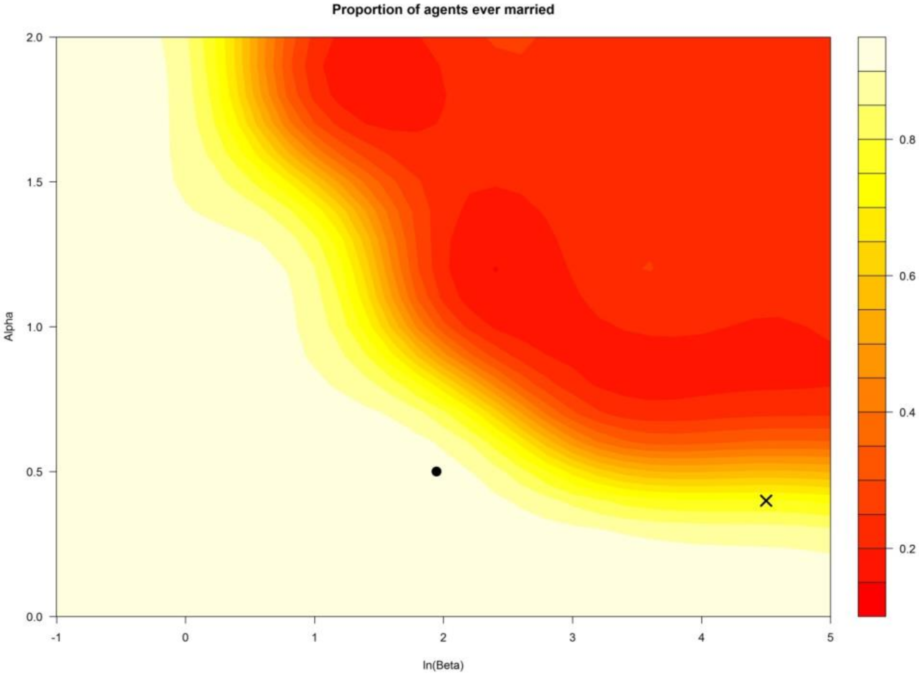

Notes: The share reported by the ONS for England and Wales in the 2011 census was 0.652 (ONS 2012a).
The black dot • shows the original parameter settings of Billari et al. (2007), and the black cross × refers to the default parameters used to generate the results reported later in the paper. The isolines are shown for $d = 25$.

Figure 3: Mean population size by social pressure parameters, simulation year 2011

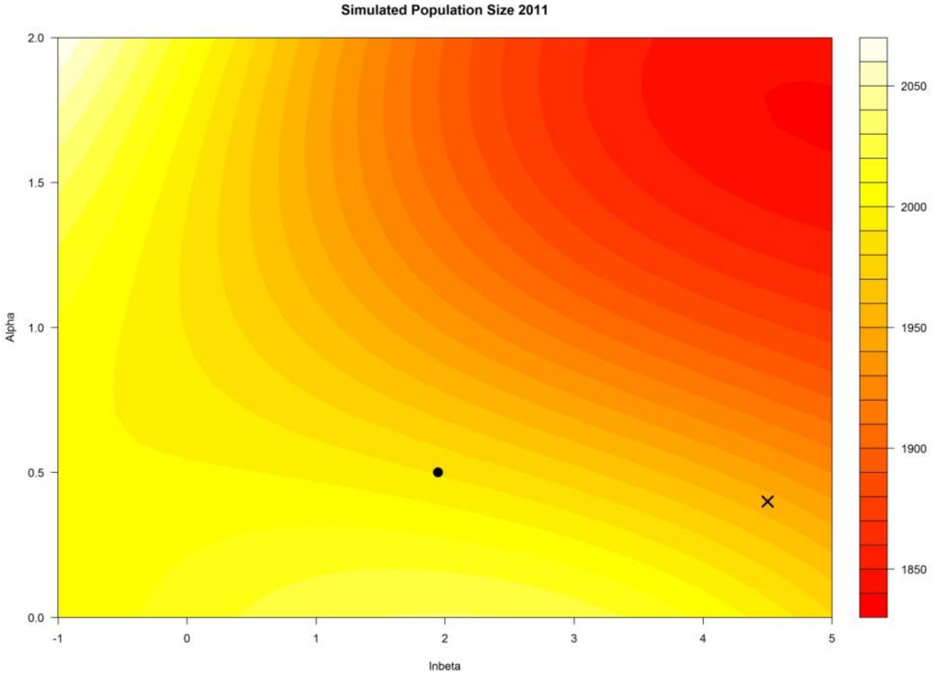

Notes: The ONS mid-2011 population estimate is 63.3 million people, corresponding to 2,014 agents; subtracting the net migration from the preceding half a century would yield 61.2 million people – 1945 agents (ONS 2013a, 2013b).
The black dot ● shows the original parameter settings of Billari et al. (2007), and the black cross × refers to the default parameters used to generate the results reported later in the paper. The isolines are shown for $d = 25$.

For population size the numbers need to be rescaled, so that the 1951 population of 50.3 million persons is represented by the 1,600 agents forming the starting population in SAMP. Proportionally, the observed mid-2011 population of 63.3 million people, adjusted upwards after the 2011 census (ONS 2013a), therefore corresponds to 2,013 agents, whereas the mid-2061 population officially projected by the ONS under the zero migration variant (63.9 million) corresponds to 2,033 agents (ONS 2012b). Knowing that cumulated net migration for the UK, since it began to be reported in 1964 until 2011, has amounted to ca. 2.1 million people (ONS 2013b), a ball-park estimate of a corresponding closed population in mid-2011 can be thus put at about 61.2 million people, that is, 1,945 agents. For the generated population size for the simulated year

2011, shown in Figure 3, the overall emulator mean of 1,959 agents corresponds to around 61.6 million people. The emulation results seem to be thus roughly in line with the empirical observations concerning both outcomes under study, the percentage of ever-married agents and overall population size, which needs to be taken into account in creating a Semi-Artificial scenario of population dynamics. The next section looks at the plausible parameterisation of such a scenario in more detail.

4.3. Illustration: A scenario with plausible marriage rates and population dynamics

In order to create a suitable Semi-Artificial scenario of population dynamics, and obtain a reasonably close match with observed empirical data, we attempt to identify a suitable set of parameters by using the emulators described in Section 4.2. These describe a system with relatively little systematic variation in population size, but with a very sharp change in the proportion of agents married with parameters α and β, as can be seen in Figure 2. When trying to interpret this outcome in demographic terms we can theorise that a form of phase transition for proportion ever-married can be observed because of the positive feedback built into the system. For small values of α or β the intensity of partner search is much more sensitive to an increased local proportion of married agents. This in turn leads to an increasing share of married agents within the network of relevant others, and thus generates a positive feedback. Conversely, when both α and β are large this critical mass is never reached, and only a small number of marriages are formed. In our system this latter case does not generally lead to declining population because of constraints that keep overall fertility close to empirical or projected rates. Still, in such cases the model would produce overall shares of more than 80%. Finally, smaller values of d result in larger populations because the criterion for the tipping point to be reached is more straightforward: only a locally high share of married agents is required. We conclude that on the space of parameters α and β there are some trade-offs between the two outputs concerned and, more generally, between different features of the underlying population dynamics.

As indicated before, a comparison with the empirical data of the UK Office for National Statistics is possible, for example, for the simulation year 2011. Given the above, then, we aim to match selected aspects of the demographic reality in 2011 by selecting parameters that the emulators indicate will give plausible outcomes for population size and proportion ever-married. Note that after all the changes were applied in our implementation of the original Wedding Ring model, the default parameter setting of Billari et al. (2007), with $\alpha = 0.5$, $\ln(\beta) = \ln(7)$, would produce overall shares of ever-married agents of around 80%: visibly higher than the empirical values. These settings are shown with dots in Figures 2 and 3. In turn, parameters α = 0.4, ln(β) = 4.5 and d = 25, depicted in Figures 2 and 3 by crosses, provide plausible outputs; this point sits on the marriage phase transition, and still retains reasonable values of the total population. For these settings we present a scenario of the dynamics of the overall population size.

The outcome of the emulation indicates that for population size the uncertainty associated with the random elements of the SAMP is nearly as significant as systematic variation. Therefore the scenario results presented in this section include a simple Monte Carlo analysis of the SAMP output, which allows for assessing the uncertainty of the model (code) itself. It has to be noted that in this case the process of matching the simulation with reality is a static exercise, which serves mainly as an illustration of the underlying concepts but could be extended to a larger number of outcomes.

Figure 4 indicates the dynamics of the simulated population over the whole period 1951–2061. Here, the mean values are shown alongside the 2.5-th and 97.5-th percentiles from the simulated set of 500 model runs. Additionally, observed population totals for 1951–2010 are presented, as well as the ones projected by the ONS for 2011–2061 in the 2010 round of National Population Projections, under the no migration assumptions. Obviously the ONS projections are benchmarked to higher values than the ones obtained in the simulation, as the latter does not take into account the positive balance of past migration into the UK. Still, the trends in the projected and simulated trajectories for the future years are very similar.

The results of this simulation indicate that the generated population trajectories are plausible from the point of view of selected empirical data and official projections. Differences between the simulated and observed trajectories are in large part due to the simplifications of SAMP, in particular the exclusion of international migration, which remains a very important component of the contemporary and projected population dynamics of the UK. Further discrepancies might result from the very basic description of the modelled marriage processes, with no explicit modelling of cohabitation, no divorce or partnership dissolution, and no re-marriage.

Figure 4: Simulated population size (black), and empirical/projected UK comparisons (red)

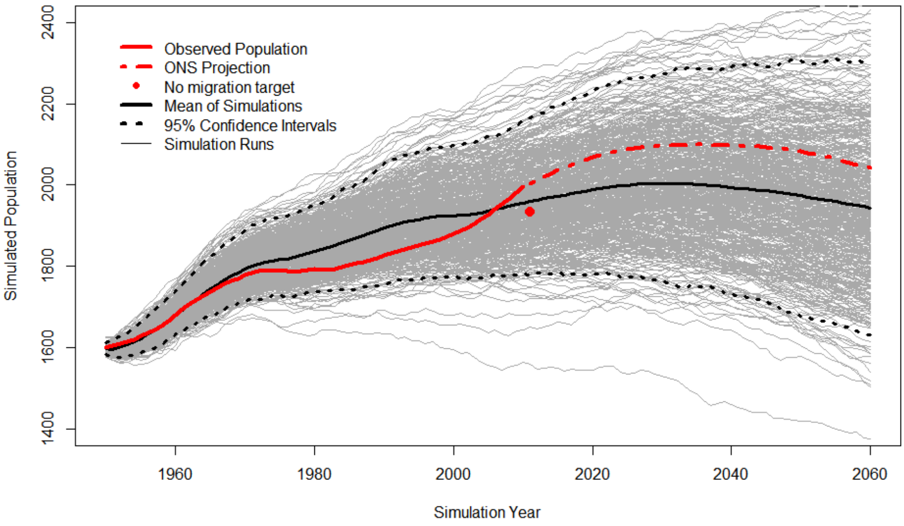

Notes: The UK population size has been rescaled so that the 1951 population matches the number of agents (1,600). The dotted lines correspond to 95% confidence bounds of the simulated population size. The observed population and the ONS projection do not take into account adjustments made after the 2011 census (ONS 2013a).
Source: (UK figures): ONS National Population Projections, 2010-based, no migration (ONS 2012b).

Figure 5 sheds greater light on the effect of some of these differences on the proportions of ever-married agents averaged across the whole simulation horizon, which are similar to patterns observed in 2011 but with slightly higher percentages married at younger ages and slightly lower for age 50 or above. In any case the effect of marriage on population is mediated by the corrections applied to fertility, as the simulated mean Total Fertility Rate remains very close to the empirical or projected values throughout the course of the simulation.

Figure 5: Empirical and simulated percentages of ever-married women by age

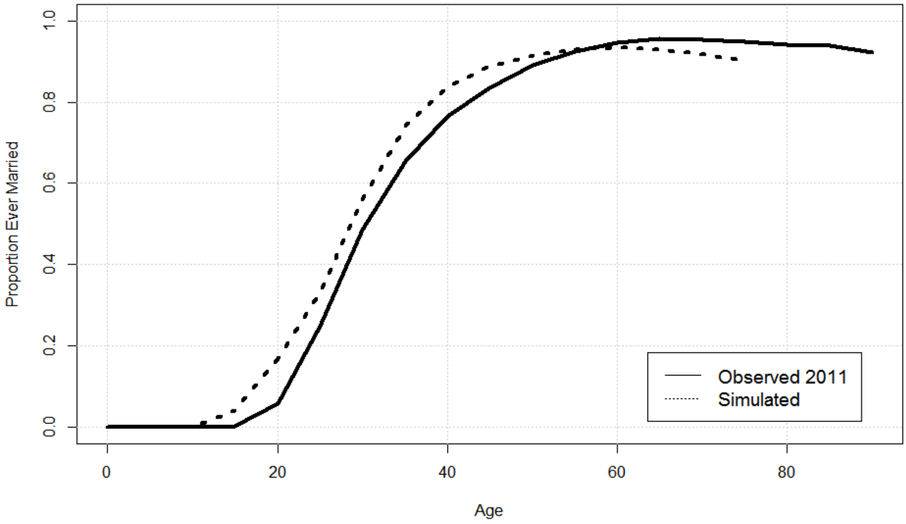

Source: Ever-married proportions for 2011, approximated by mid-2010 estimates: ONS (2011); own simulations.

One of the key advantages of agent-based models, as of all other types of micro-simulations, is that in addition to summary statistics they are able to produce whole distributions of the numbers of agents according to different characteristics. In this way the agents' heterogeneity is explicitly accounted for. One example from SAMP, showing a share of female agents from the cohort born in simulation years 1981–1990 by their completed family size, is illustrated in Figure 6. The choice of a cohort is purely illustrative, but still it demonstrates some interesting features of the simulation. Even though parity progression is not modelled explicitly in terms of agent-based rules, as is the case e.g., in the work by Aparicio Diaz et al. (2011), the share of childless female agents is very close to the one actually expected for the relevant cohort in the official population projections. Also for the higher parities (3 children and more), the projected values lie within 95% confidence bounds from the simulation. The key difference can be observed for parities 1 and 2. This indicates an area of a possible further extension of the model, should parity explicitly become an outcome of interest in future research. Such work, which would additionally need to take into account ideal family sizes and parents' preferences, remains, however, beyond the scope of the current paper.

Figure 6: Share of simulated female agents from the 1981–1990 cohort, by number of children born

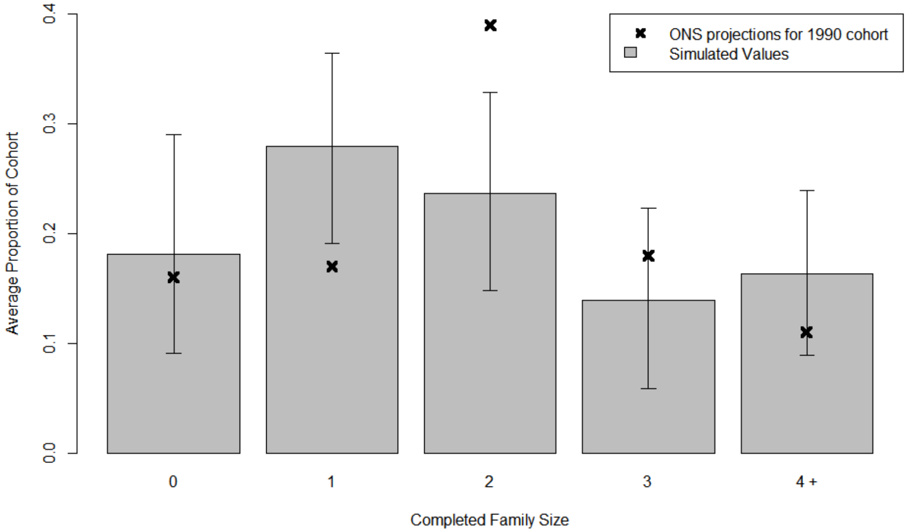

Note: Error bounds correspond to 95% confidence intervals for particular numbers of children born.
Source: Own simulations, and assumptions of the 2010-based ONS National Population Projections (ONS 2012c).

Clearly, the sample results and illustrations presented in this section are not exhaustive, since, even from such a relatively simple model as SAMP, many sets of different outputs can be produced at various levels of aggregation: from individuals, with their heterogeneity explicitly taken into account, to various aggregate population characteristics. As mentioned in the Introduction, though, the aim of this paper is mainly illustrative and exploratory rather than attempting to offer a fully realistic agent-based demography, which clearly remains an aim to strive towards. Hence, the current version of SAMP should be interpreted as a proof of concept, rather than a definitive model.

5. Conclusion

The main contribution of this paper to agent-based computational demography has been to demonstrate that using Gaussian process emulators is a convenient way of identifying plausible areas within the model parameter space, and of conducting a comprehensive analysis of uncertainty in complex computational models. In our example the sensitivity analysis indicates a key role for social pressure in the marriage formation process as implemented in the model, which proved more important than the spatial distance parameter driving the partner search. We have also shown, through a proof of concept, that agent-based models augmented with selected series of real demographic data offer improved predictive capabilities when compared to agent-based scenario generation alone. To that end, by using SAMP we have obtained the simulated population characteristics that match some of the patterns observed in the UK demography relatively well, although with some trade-offs between different output variables, as shown in Figures 2 and 3.

In general, the proposed simple multi-state model of population dynamics (SAMP) is argued to have enhanced predictive capacity as compared to the original specification of the Wedding Ring, notwithstanding the closed population assumption, and to the extent controlled by the selected inputs and outputs. Given the non-linearities in agent-based models such as the Wedding Ring, and the presence of feedback loops, the uncertainty of the model may be difficult to assess directly by using traditional statistical methods. In such cases the use of statistical emulators offers a way forward. Overall, the proposed methods allow for generating coherent, multi-level agent-based scenarios, whose increased predictive capacity is due to a combination of adding empirical basis for selected aspects of the demographic reality, and exploring the parameter space by using emulators. Emulators are also convenient for analysing the statistical properties of such models. In this way the agent-based models can be viewed through a statistical lens, reducing the gap between 'statistical' and 'simulated individuals' (cf. Courgeau 2012). To that end, we argue that these two approaches are complementary, rather than competitive.

Natural substantive extensions of models such as SAMP include other demographic processes, in particular migration (see Willekens 2012), as well as partnership dissolution and heterogeneous forms of partnerships. Fertility decisions themselves can be subject to agent-based modelling, as demonstrated by Aparicio Diaz et al. (2011), with parity distribution being an explicitly-targeted emergent outcome of the underlying behavioural rules. Another restriction of the current model consists in yearly time steps, which preclude multiple events of the same type (e.g., births) from happening in the same year, while allowing different events to occur simultaneously. In further studies these issues can be mitigated by increasing the resolution of simulation steps. Other innovations such as increasing the spatial dimensionality, which relax some of the constraints on the agents' behaviour, and adding further complexity to the state space by including the health status of agents, are reported elsewhere (Silverman et al. 2013).

Alternatively, a similar model might be tested in the context of another population; for example, where the impact of the current simplifications is envisaged to be less acute. A society with low migration levels, low rates of extra-marital childbearing, and lower divorce rates than in the UK, such as Japan, could provide a further testing ground in that respect. Another possible conceptual step might involve creating a fully artificial ('toy') model of population with data constraints lifted, where the focus would be exclusively on exploring the importance of various mechanisms driving population dynamics. Regardless of the role of data, scenario-generating models can enable policymakers to interactively examine the effects of their decisions (policy levers) under various conditions and assumptions.

Further important methodological extensions of the model would include learning about the input values from the benchmarking of outputs to the observed population characteristics, for example with respect to various summary measures of population structures, in a comprehensive manner. Such statistical calibration techniques could be explored by using full Bayesian inference in conjunction with emulators (Kennedy and O'Hagan 2001). This would allow for describing and propagating scenario uncertainty stemming from different sources, not only the model code, in a coherent way. In particular this approach could be applied to calibrating the emulator results against the series of historical data, in a process known as history matching[12]. This is an especially promising avenue of further research, given the dynamic nature of the system under study. Finally, more work should be done on the design of the experimental space, $\mathbf{\Delta}$, for example by using Latin Hypercube samples or randomisation (O'Hagan 2006).

The substantial and methodological challenges notwithstanding, given that population is still one of the better-predictable social processes despite the surrounding uncertainty, demography can offer a testing ground for social science applications of emulator-based methods for dealing with uncertainty in complex computational models. If successful, this approach can be a means of addressing two main methodological challenges of contemporary demography: the inclusion of micro-level mechanistic explanations of patterns observed at the population level, and the coherent description of uncertainty at different levels of aggregation.

---

[12] See: http://mucm.aston.ac.uk/MUCM/MUCMToolkit/index.php?page=ThreadGenericHistoryMatching.html (accessed on 20 May 2013).

6. Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) grant EP/H021698/1 “Care Life Cycle”, funded within the ‘Complexity Science in the Real World’ theme. The paper offers a demographic extension of the prototype model reported in Silverman et al. (2013). The authors are very grateful to Daniel Courgeau and Frans Willekens for inspiring methodological discussions on agent-based demography; as well as to two anonymous Reviewers and the Associate Editor of *Demographic Research*, Alexia Fürnkranz-Prskawetz, for really helpful comments that helped us improve the earlier draft. All the remaining errors are exclusively ours, and we retain the sole responsibility for all views and interpretations reported in this paper.

References

Aparicio Diaz, B., Fent, T., Prskawetz, A., and Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography* 48(2): 559–579. doi:10.1007/s13524-011-0023-6.

Axtell, R.L., Epstein, J.M., Dean, J.S., Gumerman, G.J., Swedlund, A.C., Harburger, J., Chakravarty, S., Hammond, R., Parker, J., and Parker M. (2002). Population growth and collapse in a Multi-Agent Model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciencesof the USA* 99(Suppl. 3): 7275–7279. doi:10.1073/pnas.092080799.

Benenson, I., Omer, I., and Hatna, E. (2003). Agent-Based modeling of householders' migration behaviour and its consequences. In: Billari, F. and Prskawetz, A. (eds.). *Agent–Based Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behaviour*. Heidelberg: Physica Verlag: 97–115. doi:10.1007/978-3-7908-2715-6_6.

Bernardi, L. (2003). Channels of social influence on reproduction. *Population Research and Policy Review* 22(5–6): 527–555. doi:10.1023/B:POPU.0000020892.15221.44.

Bernardi, L., Keim, S., and von der Lippe, H. (2007). Social influences on fertility: A comparative mixed methods study in Eastern and Western Germany. *Journal of Mixed Methods Research* 1(1): 23–47. doi:10.1177/2345678906292238.

Billari, F., Aparicio Diaz, B., Fent, T., and Prskawetz, A. (2007). The "Wedding–Ring": An agent–based marriage model based on social interaction. *Demographic Research* 17(3): 59–82. doi:10.4054/DemRes.2007.17.3.

Billari, F., Fent, T., Prskawetz, A., and Scheffran, J. (2006). *Agent–Based computational modelling: Applications in demography, social, economic and environmental sciences*. Heidelberg: Physica Verlag.

Billari, F. and Prskawetz, A. (2003). *Agent–Based computational demography: Using simulation to improve our understanding of demographic behaviour*. Heidelberg: Physica Verlag. doi:10.1007/978-3-7908-2715-6.

Billari, F.C. (2006). Bridging the gap between micro-demography and macro-demography. In: Caselli, G., Vallin, J., and Wunsch, G. (eds.). *Demography: analysis and synthesis, Vol 4*. New York: Academic Press / Elsevier: 695–707.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the USA* 99(Suppl. 3): 7280–7287. doi:10.1073/pnas.082080899.

Booth, H., Maindonald, J., and Smith, L. (2002). Applying Lee-Carter under conditions of variable mortality decline. *Population Studies* 56(3): 325–336. doi:10.1080/00324720215935.

Bracher, M., Watkins, S., and Santow, G. (2003). 'Moving' and Marrying: Modelling HIV Infection among Newly-weds in Malawi. *Demographic Research, Special Collection* 1(7): 207–246. doi:10.4054/DemRes.2003.S1.7.

Bühler, Ch. and Frątczak, E. (2007). Learning from others and receiving support: The impact of personal networks on fertility intentions in Poland. *European Societies* 9(3): 359–382. doi:10.1080/14616690701314101.

Courgeau, D. (2012). *Probability and social science: Methodological relationships between the two approaches. Methodos Series, vol. 10.* Dordrecht: Springer.

Dannefer, D. (2003). Cumulative advantage/disadvantage and the life course: Cross-fertilizing age and social science theory. *Journal of Gerontology (Social Sciences)* 58(16): S327–S337. doi:10.1093/geronb/58.6.S327.

Entwisle, B. (2007). Putting people into place. *Demography* 44(4): 687–703. doi:10.1353/dem.2007.0045.

Epstein, J.M. (2008). Why model? *Journal of Artificial Societies and Social Simulation* 11(4): 12. http://jasss.soc.surrey.ac.uk/11/4/12.html.

Epstein, J.M. and Axtell, R.L. (1996). *Growing artificial societies: Social science from the bottom up*. Cambridge, MA: MIT Press.

Eurostat (2011). Eurostat Statistics Database; Domain Population and Social Conditions [electronic resource]. Luxembourg: Eurostat. http://epp.eurostat.ec.europa.eu (data retrieved on 27/10/2011).

Geard, N., McCaw, J.M., Dorin, A., Korb, K.B., and McVernon, J. (2013) Synthetic population dynamics: A model of household demography. *Journal of Artificial Societies and Social Simulation* 16(1): 8. http://jasss.soc.surrey.ac.uk/16/1/8.html.

Gilbert, N. and Tierna, P. (2000). How to build and use agent-based models in social science. *Mind and Society* 1(1): 57–72. doi:10.1007/BF02512229.

Gilbert, N. and Troitzsch, K.G. (2005). *Simulation for the Social Scientist. 2nd ed*. Maidenhead: Open University Press.

Hammel, E.A., McDaniel, C.K., and Wachter, K.W. (1979). Demographic consequences of incest tabus: A microsimulation analysis. *Science* 205(4410): 972–977. doi:10.1126/science.205.4410.972.

Heiland, F. (2003). The collapse of the Berlin Wall: Simulating state-level East to West German migration patterns. In: Billari, F. and Prskawetz, A. (eds.). *Agent–Based computational demography: Using simulation to improve our understanding of demographic behavior*. Heidelberg: Physica-Verlag: 73–96.

Hills, T. and Todd, P. (2008). Population heterogeneity and individual differences in an assortative Agent-Based Marriage and Divorce Model (MADAM) using search with relaxing expectations. *Journal of Artificial Societies and Social Simulation* 11(4): 5. http://jasss.soc.surrey.ac.uk/11/4/5.html.

Human Mortality Database (2011). Human Mortality Database [electronic resource]. Berkeley, CA, Rostock: University of California, Berkeley, Max Planck Institute for Demographic Research. [http://www.mortality.org] (data retrieved on 26/07/2011).

Hutchins, E. and Hazelhurst, B. (1995). How to invent a lexicon: the development of shared symbols in interaction. In: Gilbert, N. and Conte, R. (eds.). *Artificial societies: The computer simulation of social life*. London: University College London Press: 157–189.

Hyndman, R. (2011). forecast: Forecasting functions for time series [electronic resource]. Melbourne: Monash University. [http://robjhyndman.com/software/forecast/] (retrieved on 06/09/2011).

Keilman, N. and Pham, D.Q. (2006). *Prediction intervals for Lee-Carter-based mortality forecasts*. Paper presented at the European Population Conference, Liverpool, UK, 21–24 June 2006. http://epc2006.princeton.edu/download.aspx?submissionId=60211.

Kennedy, M. (2004). Description of the Gaussian process model used in GEM-SA. Software manual [electronic resource]. Sheffield: Centre for Terrestrial Carbon Dynamics, University of Sheffield. [http://ctcd.group.shef.ac.uk/gem.html] (retrieved on 23/05/2012).

Kennedy, M.C. and O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society, Series B* 63(3): 425–464. doi:10.1111/1467-9868.00294.

Kniveton, D., Smith, Ch., and Wood, S. (2011). Agent-based model simulations of future changes in migration flows for Burkina Faso. *Global Environmental Change* 21(S1): S34–S40. doi:10.1016/j.gloenvcha.2011.09.006.

Langton, C.G. (1995). *Artificial life: An overview*. Cambridge, MA: MIT Press.

Lee, R.D. and Carter, L.R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87(419): 659–671. doi:10.2307/2290201.

Lee, R.D. and Tuljapurkar, S. (1994). Stochastic population forecasts for the United States: Beyond high, medium, and low. *Journal of the American Statistical Association* 89(428): 1175–1189. doi:10.1080/01621459.1994.10476857.

Murphy, M. (2003). Bringing behavior back into micro-simulation: Feedback mechanisms in demographic models. In: Billari, F. and Prskawetz, A. (eds.). *Agent–Based computational demography: Using simulation to improve our understanding of demographic behavior.* Heidelberg: Physica-Verlag: 159–174.

Murphy, M. (2004). Tracing very long-term kinship networks using SOCSIM. *Demographic Research* 10(7): 171–196. doi:10.4054/DemRes.2004.10.7.

Oakley, J. and O’Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika* 89(4): 769–784. doi:10.1093/biomet/89.4.769.

Oakley, J. and O’Hagan, A. (2004). Probabilistic sensitivity analysis of complex models: A Bayesian approach. *Journal of the Royal Statistical Society, Series B* 66(3): 751–769. doi:10.1111/j.1467-9868.2004.05304.x.

ONS (1998). Birth Statistics. Series FM1, No. 27 [electronic resource]. London: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/vsob1/birth-statistics--england-and-wales--series-fm1-/no--27--1998/birth-statistics-series-fm1.pdf] (retrieved on 15/11/2011).

ONS (2011). Marital Status Population Estimates Excluding Marriages Abroad, Mid 2010 [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/pop-estimate/population-estimates-by-marital-status/mid-2010/rft---marital-status-estimates-excluding-marriages-abroad-mid-2010.zip] (retrieved on 15/08/2013).

ONS (2012a). 2011 Census: Key Statistics for England and Wales, March 2011 [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/census/2011-census/key-statistics-for-local-authorities-in-england-and-wales/stb-2011-census-key-statistics-for-england-and-wales.html#tab---Marital-status] (retrieved on 15/08/2013).

ONS (2012b). Table J1-1, Natural change variant. UK summary, 2010-based [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/npp/national-population-projections/2010-based-projections/rft-table-j1-1-natural-change-variant---uk-summary.xls] (retrieved on 15/08/2013).

ONS (2012c). Chapter 3: Fertility, 2010-based NPP Reference Volume [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/npp/national-population-projections/2010-based-reference-volume--series-pp2/fertility.html] (retrieved on 15/08/2013).

ONS (2013a). Annual Mid-year Population Estimates, 2011 and 2012 [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/rel/pop-estimate/population-estimates-for-uk--england-and-wales--scotland-and-northern-ireland/mid-2011-and-mid-2012/stb---mid-2011---mid-2012-uk-population-estimates.html] (retrieved on 15/08/2013).

ONS (2013b). Long-term Migration into and out of the United Kingdom: Timeline [electronic resource]. Titchfield: Office for National Statistics. [http://www.ons.gov.uk/ons/interactive/theme-pages-1-1/index.html] (retrieved on 15/08/2013).

Orcutt, G., Caldwell S., and Wertheimer II, R. (1976). *Policy exploration through microanalytic simulation*. Washington DC: Urban Institute.

Orcutt, G.H. (1957). A new type of socio-economic system. *The Review of Economics and Statistics* 39(2): 116–123. doi:10.2307/1928528.

O'Hagan, A. (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety* 91(10–11): 1290–1300. doi:10.1016/j.ress.2005.11.025.

Reeves, J.H. (1987). Projection of Number of Kin. In: Bongaarts, J., Burch, T.K., and Wachter, K.W. (eds.). *Family demography: Methods and their application*. Oxford: Clarendon Press: 228–248.

Ruggles, S. (1987). *Prolonged connections: The rise of the extended family in nineteenth century England and America*. Madison, WI: University of Wisconsin Press.

Ruggles, S. and De Quincey, T. (1993). Confessions of a Microsimulator : Problems in modeling the demography of kinship. *Historical Methods* 26(4): 161–169. doi:10.1080/01615440.1993.9956352.

Schelling, T.C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology* 1(2): 143–186. doi:10.1080/0022250X.1971.9989794.

Schelling, T.C. (1978). *Micromotives and macrobehavior*. Norton, New York, USA.

Silverman, E., Bijak, J., Hilton, J., Cao, V.D., and Noble, J. (2013, forthcoming) When demography met social simulation: A tale of two modelling approaches. *Journal of Artificial Societies and Social Simulation*.

Silverman, E., Bijak, J., and Noble, J. (2011) Feeding the beast: Can computational demographic models free us from the tyranny of data? In: Lenaerts, T., Giacobini, M., Bersini, H., Bourgine, P., Dorigo, M., and Doursat, R. (eds.). *Advances in Artificial Life, ECAL 2011*. Cambridge, MA: MIT Press: 747–754.

Silverman, E. and Bryden, J. (2007) From artificial societies to new social science theory. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., and Coutinho, A. (eds.). *Advances in artificial life, 9th European Conference, ECAL 2007*. Berlin-Heidelberg: Springer: 565–574.

Smith, J.E. (1987). The computer simulation of kin sets and kin counts. In: Bongaarts, J., Burch, T.K., and Wachter, K.W. (eds.). *Family demography: Methods and their application*. Oxford: Clarendon Press: 249–266.

Spielauer, M. (2007). Dynamic microsimulation of health care demand, health care finance and the economic impact of health behaviours: Survey and review. *International Journal of Microsimulation* 1(1): 35–53.

Thomson, E., Winkler-Dworak, M., Spielauer, M., and Prskawetz, A. (2012). Union instability as an engine of fertility? A microsimulation model for France. *Demography* 49(1): 175–195. doi:10.1007/s13524-011-0085-5.

Todd, P.M., Billari, F.C., and Simão, J. (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography* 42(3): 559–574. doi:10.1353/dem.2005.0027.

van Imhoff, E. and Post, W. (1998). Microsimulation methods for population projection. *Population. English selection* 10(1): 97–138.

Wachter, K.W. (1987). Microsimulation of household cycles. In: Bongaarts, J., Burch, T.K., and Wachter, K.W. (eds.). *Family demography: Methods and their application*. Oxford: Clarendon Press: 215–227.

Wachter, K.W. (1995). *2030's Seniors: Kin and Step-Kin [unpublished manuscript]*. Berkeley, CA: University of California.

Wachter, K.W. (1997). Kinship resources for the elderly. *Philosophical transactions of the Royal Society of London. Series B: Biological sciences* 352(1363): 1811–1817. doi:10.1098/rstb.1997.0166.

Willekens, F. (2005). Biographic forecasting: Bridging the micro-macro gap in population forecasting. *New Zealand Population Review* 31(1): 77–124.

Willekens, F. (2012). *Migration: A perspective from complexity science*. Paper presented at the Migration workshop of the Complexity Science for the Real World network, Chilworth, UK, 16 February 2012.

Wolf, D. (1994). The elderly and their kin: Patterns of availability and access. In: Martin, L. and Preston, S. (eds.). *Demography of aging*. Washington DC: National Academies Press: 146–195.

Zinn, S., Gampe, J., Himmelspach, J., and Uhrmacher, A.M. (2009). MIC–CORE: A tool for microsimulation. In Rossetti, M.D., Hill, R.R., Johansson, B., Dunkin, A., and Ingalls, R.G. (eds.). *Proceedings of the 2009 Winter Simulation Conference*. Austin, TX: IEEE: 992–1002. doi:10.1109/WSC.2009.5429424.