

Semi-Artificial Models of Populations: Connecting Demography with Agent-Based Modelling

Eric Silverman*, Jakub Bijak, Jason Noble, Viet Cao, and
Jason Hilton

University of Southampton,
Southampton, SO17 1BJ, United Kingdom
{e.silverman,j.bijak,jn2,vdc1g10,jdh4g10}@soton.ac.uk
<http://www.southampton.ac.uk/clc>

Abstract. In this paper we present an agent-based model of the dynamics of mortality, fertility, and partnership formation in a closed population. One of our goals is to bridge the methodological and conceptual gaps that remain between demography and agent-based social simulation approaches. Model construction incorporates elements of both perspectives, with demography contributing empirical data on population dynamics, subsequently embedded in an agent-based model situated on a 2D grid space. While taking inspiration from previous work applying agent-based simulation methodologies to demography, we extend this basic concept to a complete model of population change, which includes spatial elements as well as additional agent properties. Given the connection to empirical work based on demographic data for the United Kingdom, this model allows us to analyse population dynamics on several levels, from the individual, to the household, and to the whole simulated population. We propose that such an approach bolsters the strength of demographic analysis, adding additional explanatory power.

Keywords: Demography, Agent-Based Modelling, Social Simulation

1 Motivations: Bringing together the statistical and the simulated individual

Silverman and Bryden [22] divided current work in agent-based models (ABMs) in social science into two main streams: systems sociology and social simulation. The former focuses on agent-based models in an explanatory role, with few predefined interactions or structures; the models exist to probe elements of social theory, rather than to examine the functionality of specific social systems (e.g., investigations of the behaviour of nation-states, as in [8]; or the social effects of different kin systems, [19]). Social simulation approaches instead focus on a

* This work was supported by the EPSRC grant EP/H021698/1 Care Life Cycle, funded within the Complexity Science in the Real World theme.

particular class of behaviour or a particular social system, and may have some link to empirical data (e.g., the reconstruction of the Anasazi population in [1]; or [11] on the movement of early Polynesian peoples).

In general, agent-based models have recently become increasingly popular in the social sciences, as the methodology has become more established. The models have produced useful insights even when no empirical data was used at all (e.g., [20]), and therefore seem to have demonstrated a certain capacity for explanatory power. Meanwhile, some researchers have proposed that ABMs will allow social scientists to more directly integrate their data with simulation approaches [18]. Given the remarkable flexibility the ABMs display, such models can be used for numerous purposes beyond simply explanation and prediction [12].

The methodological challenges facing demography currently (see [21]) have led some demographers to investigate the methodology of agent-based modelling as a potentially valuable addition to the toolkit of the field [6], [4]. As mentioned before, specifically in the context of demography, the challenges facing multi-level microsimulation models have alerted some authors to the need to address the particular issue of the combinatorial explosion of the parameter space in these models [21]. In addition, it has been argued that traditional demographic models cannot fully capture the complexities of micro-level agent behaviour and heterogeneity, nor can they address downward causation from macro-level entities or groups (*idem*). Agent-based models can provide these possibilities [13], as well as a possible platform for understanding social interactions, social networks, and other processes lying at the root of demographic change [6].

Recent work in demography has described the fundamental goals of demographic research – or indeed any field of social science applying statistical methods – as the use of statistical models based on observations (e.g., surveys, censuses) to predict and describe the behaviour of statistical individuals [10]. Demography has therefore progressed through a number of different paradigms, ranging from the cross-sectional view, in which events are viewed as separate from individuals, to the event-history perspective, in which demographic events are interdependent on individual life-courses. In recent years, the most prominent view has been the multilevel paradigm, in which behaviour at the societal level can only be understood by investigating the multiple levels of the social world simultaneously ([9], [10]).

A similar process can be seen within the systems sociology and social simulation communities: in these arenas we focus on simulated individuals equipped with some simple behavioural rules, through which we hope to see the emergence of interesting behaviours at the macro level. At the same time, some of these models incorporate feedback loops, downward causation effects, social networks and other related concepts – all of which serve to bring these multiple levels of the social into play within the simulation.

Thus, we argue that while social science ABMs and demography clearly differ in some respects – for example having either theoretical or empirical focus, using mainly computational or statistical models, etc. – in fact the recent his-

tory of demographic theory shows synergies developing between the agent-based simulation approach and the multilevel demographic approach. On the technical side, we already see some of these links: some multi-state models in demography already assume simple behavioural rules in individuals, while some ABMs of social systems already use substantial empirical information in the model-building process.

Moreover, demographic methods also have significant strengths in the empirical arena, not least the capacity for higher predictive accuracy in many cases, as well as the direct connection to the rich information already embedded in the age structure of populations. Efforts thus far in developing agent-based demography have attempted to incorporate agent-based methods while simultaneously retaining these particular strengths of demography (e.g., [6]; [15]; [23]; [4]; [5]). Thus, demographic methods can contribute significant additional expertise and allow us to better align models with the real world of empirical observations.

Taking this process further, we propose that directly linking demographic methods with ABM frameworks will allow us to produce models which increase our understanding of population change – while simultaneously helping us to avoid the pitfalls of an overdependence on empirical data [21]. ABMs allow us to produce models which have a greater explanatory capacity, while the demographic components allow to use the inherent flexibility of the ABM approach to generate plausible scenarios within a given parameter space, informed by well-established demographic methods of analysis.

In this framework, such augmented ABMs will allow demographers to examine scenarios over longer time horizons, rather than being heavily limited by data-dependent statistical models. The general consensus is that the predictive horizon of demographic models is about one generation in length [16], but some have argued that scenario-based methods can offer an opportunity to explore a larger space of possible futures (e.g., [24]; [3]). Thus, we propose that this new framework can produce a shift within the demographic community: from a focus on prediction and description alone, to an approach which allows for exploration of scenarios of population change, while providing some insight into the underlying processes.

Hence, we suggest that the social simulation stream of agent-based modelling is of great relevance to demography. The model presented here aims to make progress at this interface, and bring agent-based models of population to the fore as a means for both enhancing demographic methods, and for tying agent-based models more closely to real social systems (an approach we have advocated previously; [21]). We refer to this approach as a semi-artificial model of a population to capture the mixture between empirical data and randomly-generated agent populations with simple behavioural rules. An example of such a model – inspired by and drawing from the Wedding Ring model seen in [5] – is described in detail in the next section.

2 From the Wedding Ring to the Wedding Doughnut

2.1 Wedding Ring: Background

In order to illustrate the potential benefits of combining demographic and complexity science approaches, we replicated and expanded upon the model of partnership formation designed by Billari et al. [5], known as the Wedding Ring model. This model attempts to explain the age-at-marriage patterns seen in modern European states by representing the process of partnership formation as a consequence of social pressure. This 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., [2]; [7]; [5]).

Thus, the model represents the spread of marriage through a population as a diffusion process. However, marriage differs from other diffusion processes in that even those experiencing a very high level of social pressure towards marriage still cannot get married without finding a suitable unmarried partner [5].

The Wedding Ring is so named due to the fact that agents live in a one-dimensional ring-shaped world, which changes over time (*idem*). The agents are thus effectively situated in a cylindrical space, with the circular dimension (arc length) representing space, and the linear dimension – time. Each agents 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. The overall level of social pressure and the agents age influence parameter determine the range in which agents search for a suitable partner. 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.

In order to define the network of relevant others in which the agent searches for a partner, each agent is first classified 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). The size of the spatial interval for the agents network of relevant others is symmetric around their location, and varies according to the size of the initial population.

2.2 Extending the Model: the Wedding Doughnut

In order to better align our model with the demographic processes under study, we altered the Wedding Ring model in several major ways. First, we situated the

agents in a toroidal space, as this would allow for a more complete consideration of the impact of the spatiality aspect in this implementation of partnership formation. As in the original Wedding Ring, the dimensions of the grid space could be considered not just as spatial distances, but also social distances [5], or, in our example, a combination of both. The model uses a grid space 72 squares long on each edge, with an initial population of 800 agents – this allowed for a sufficient population density to observe interesting population dynamics.

In our model – hereafter dubbed the Wedding Doughnut, given the new spatial arrangements – several core components of the original model were altered. The Wedding Ring was built on the assumption that each agent only had one spatial coordinate to record, so in the Doughnut we altered the methods used to calculate spatial separation, and substantially altered other parameters in order to allow the agents to properly search this space.

We also allowed the agents to move on the grid. When an agent forms a partnership, they move to a new location, the distance to which is inversely proportional to the size of the network of their relevant others. Any future offspring are placed at this new household. For simplicity, however, we follow Billari et al.s lead and assume that partnerships are permanent, and that agents cannot form partnerships until the age of 16. Future extensions of this model might allow for more detail, such as same-sex partnerships and the possibility of partnership dissolution.

The other details of the model have not changed in this implementation – we use an identical function to the Wedding Ring for age influence (a piecewise linear function which varies with age) and for social pressure (a sigmoid function). We used identical parameter values for these as in Billari et al. [5] in their base scenario.

2.3 Extending the Model: Demographic Elements

The original Wedding Ring made some significant simplifying assumptions in order to make the model run smoothly – agents died only when reaching age 100, and birth rates were adjusted regularly in order to keep the population constant. Both of these restrictions were removed in the Doughnut model.

Further, we wished to include in this extension a serious consideration of the underlying demographic processes that influence partnership formation. Modern societies are ageing, and older agents have smaller networks of relevant others – therefore we added a realistic model of mortality, as otherwise the model risks failing to capture the complexities underlying these trends. Similarly, we added a realistic model of fertility to capture the current shift toward later child-bearing and lower birth-rates.

The initial population in the model is generated at random, but the distributions of agents by age, sex and marital status correspond to the observed data from England and Wales in the 1951 Census. As the simulation progresses, fertility and mortality rates are based on empirical data and projections for the UK population. The first 60 years of the simulation use age-specific mortality rates drawn from the Human Mortality Database (2011) for 1951–2009.

These rates were then projected forward using the popular method for forecasting mortality developed by Lee and Carter [17]. This method uses the leading vectors of a singular value decomposition of the matrix of centred mortality rates to construct a model for mortality with only one time varying element. This allows easy forecasting using standard times series methods; more details about procedure and estimation are available in [17]. Projections to 2250 using this method show a continual but slowing increase in life expectancy over the period.

Fertility rates for the simulation were obtained in a similar way. Age-specific fertility rates from 1973–2009 for UK women of childbearing age were obtained from the Eurostat database (Eurostat 2011), while earlier data for the period 1950–1972 were taken from the Office of National Statistics data for England and Wales.

A Lee-Carter model was again fitted to the data to obtain future rates, but, in contrast to the mortality projections, two components of the singular value decomposition of the matrix of fertility rates were needed, as two time indices better captured the trends in fertility. The resultant projections to 2250 sees the total fertility rate increase initially before converging at a value just above replacement fertility, and also display a continuation in the empirical trend towards later childbearing in the UK.

2.4 Extending the Model: Health Status

In order to make the model more relevant to real-world policy concerns related to the ageing UK population, particularly problems in social care provision, we also added a simplistic model of health status to the Wedding Doughnut. Agents have a probability of transitioning into a state of ill health, which increases with age. Once agents transition into this state, they remain ill until death – we thus tacitly assume that this state indicates the type of limiting long-term illness that requires significant care provision. As in the real world, males have a slightly higher probability of entering this state than females; parameter values were not based on empirical data, due to the difficulty of obtaining and analysing data on these types of illnesses.

While this model of health status is extremely simplistic, we intended this addition to serve as a proof-of-concept that a simulation built on this semi-artificial framework could produce results which, due to their relationship to empirical data and projections, could bear upon issues relevant to policy-makers. As the results will show, running various scenarios even in this basic model produced interesting results – and future iterations of the model could additionally incorporate a more robust model of health status which would allow for significantly more detail.

2.5 The Wedding Doughnut: Flow of the Simulation

The Wedding Doughnut was written in Repast Symphony v. 2.0, which is a Java-based software package for agent-based modelling. Figure 1 provides an illustration of the general flow and structure of the model. The simulation runs

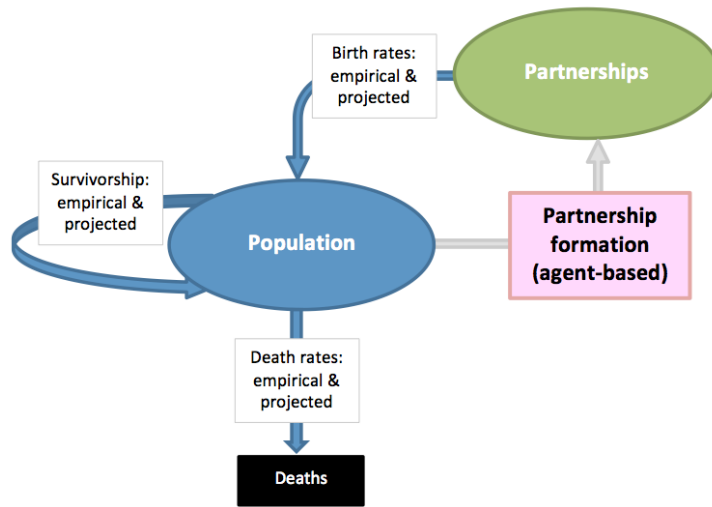


Fig. 1. A representation of the flow of the Wedding Doughnut model. Agent-based methods are used to produce a model of partnership formation, which is linked to the other empirical, demographic methods shown here to produce the complete model.

for 300 time-steps, each time-step being one year - this is in contrast to the Wedding Ring, which only ran for 150 years. The first time-step corresponds to the calendar year 1951, and hence the simulation extends to 2250.

Each time-step proceeds as follows:

1. All agents age one year
2. For agents without partners:
 - (a) For each agent, we find their relevant others
 - (b) Social pressure is calculated
 - (c) Potential partners are found
 - (d) Partnerships form where there is mutual need/pressure
 - (e) New partners move to a new location
3. For agents with partners:
 - (a) Check fertility status
 - (b) Some agents will give birth according to relevant fertility rates
4. For every agent:
 - (a) Check health status: Some agents will become ill according to relevant rates
 - (b) Check mortality: Some agents will die according to relevant mortality rates
5. Remove dead agents from the simulation
6. Place newborn agents into the simulation

3 Results

Initial runs of the model were designed as a pure replication of the original Wedding Ring: agents were placed on a 1D ring, and none of our modifications were incorporated. Results were consistent with the original paper, which indicated that we could progress to developing our modifications.

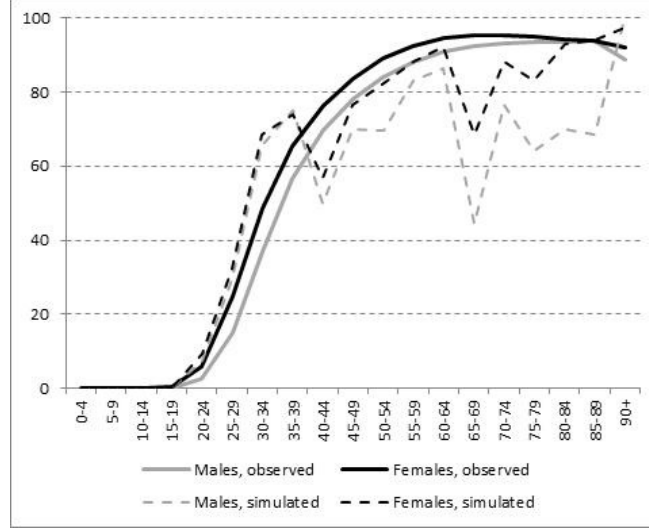


Fig. 2. Example of a comparison between observed and simulated populations showing the percentage of the population which has ever married by 2010–2011.

Once our modifications based on empirical demographic projections were included, our initial set of 1800 runs indicated qualitative similarity to the patterns of marriage observed in modern Britain. As shown in Figure 2, the simulation produced populations with proportions of individuals who married at some point during their life-course as we see during 2010–11 in the UK. The results show much greater stochasticity, which is likely due to our small initial population, and the short running time of the simulation by the time it reaches calendar year 2011.

Using our initial parameter settings, we were able to broadly replicate the pattern of marriage that we see in empirical data. Sensitivity analysis further showed that significantly altering the basis of the simulation by reducing social pressure or age influence parameters to constants produced marriage hazard rates inconsistent with reality (as shown in Figure 3). The base scenario, in which age influence varies according to the age of the agents and social pressure is calculated using a sigmoid function, more closely resembles hazard rates of marriage that we see in modern European states.

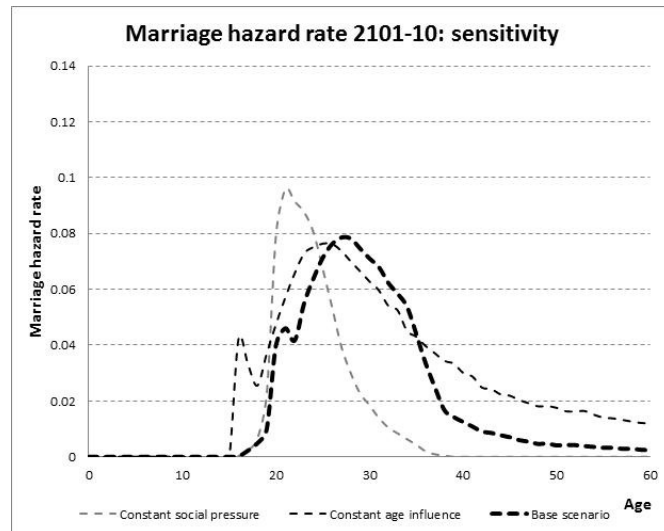


Fig. 3. Comparison of marriage hazard rate results in three scenarios: the base scenario; constant social pressure; and constant age influence.

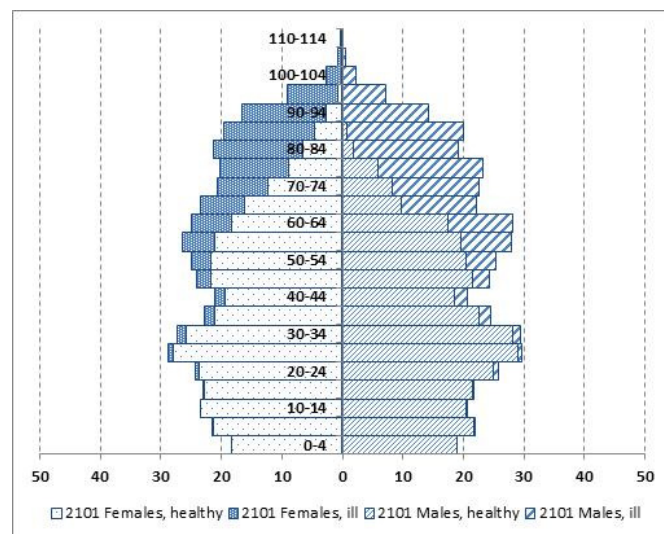


Fig. 4. Population pyramid for the year 2101 in our simulated population. Note the increasing incidence of long-term limiting illness as age increases.

Figures 4 and 5 provide a breakdown of a simulated population in the year 2101. While these results obviously cannot be compared with empirical data, we do see a pattern here that intuitively fits our expectations for future population dynamics. In Figure 4, we see that the population becomes increasingly dominated by the older age brackets of society, and further, many of those individuals have taken ill and require significant care. Figure 5 shows that as our simulated individuals age, many of those who are ill have not had the opportunity to marry.



Fig. 5. Breakdown of simulated population in the year 2101 by both marriage and health status.

Thus, results show an encouraging parity with empirical data in the early stages of the simulation (which are the only stages for which we can make this comparison). Hazard rates of marriage reflect what we expect in current society, and overall marriage rates, while displaying more stochasticity, are broadly at a level consistent with reality. As the simulation progresses, we see population change that mirrors the expectations of demographers with regard to the dynamics of an ageing population. Agent populations become increasingly dominated by the old and infirm, marriage/fertility rates slowly decline, and as a consequence we see ever-increasing numbers of ill agents who require care but who have no family which could provide it.

4 Discussion

The results above demonstrate our original thesis: that agent-based models of demographic processes, augmented with empirical data on fertility and mortality, can produce results that match our expectations of the demographics of

modern European states. Despite the lack of data fully informing our model of health status, we see plausible distributions of healthy and unhealthy agents that illustrate the consequences of Europe's ageing population. This leads us to conclude that the mechanisms used here to drive basic demographic processes have captured the essential elements needed to produce a useful starting point for semi-artificial population models for the United Kingdom.

While this is a promising start, a number of further extensions are required. First, our model of health status, while illustrative, is quite simplistic and does not necessarily reflect the complexity of relationships between health and ageing in the modern UK society. As such, this part of the model can be potentially extended by incorporating empirical data on limiting long-term illness in the UK. Second, our model of partnership seems to capture the appropriate dynamics, but it does not incorporate the possibility of partnership dissolution, or of same-sex partnerships. Finally, agents on our Wedding Doughnut do not shift position in that space unless they have formed a new partnership. Natural extensions in that respect would require incorporating more sophisticated agent behaviours, with particular interest toward allowing for agent mobility that allow us to represent the possible causes and effects of within-country migration.

Nevertheless, even in its current form, the Wedding Doughnut provides an illustrative example of the power of the combined ABM-demography approach. We propose that this approach can help both disciplines in combination to improve their explanatory relevance, deepen our understanding of demographic processes, and increase our appreciation of the links between macro-level effects and micro-level behaviour. This may move us closer to the ideal expressed by the British demographer John Hajnal, who aspired to the construction of models which involve [...] more cognition than has generally been applied ([14], p. 321).

References

1. Axtell, R.L., Epstein, J.M., Dean, J.S., Gumerman, G.J., Swedlund, A.C., Harberger, J., Chakravarty, S., Hammond, R., Parker, J. and Parker, M. (2002). Population Growth and Collapse in a MultiAgent Model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences, Colloquium* 99(3): 7275–7279.
2. Bernardi, L. (2003). Channels of social influence on reproduction, *Population Research and Policy Review*, 22(56): 527–555.
3. Bijak, J. (2010). *Forecasting International Migration in Europe: A Bayesian View*. Springer Series on Demographic Methods and Population Analysis, vol. 24. Dordrecht: Springer.
4. Billari, F.C. (2006). Bridging the gap between micro-demography and macro-demography. In: G. Graziella, J. Vallin and G. Wunsch (eds.) *Demography: analysis and synthesis*, Vol 4. New York: Academic Press / Elsevier (pp. 695–707).
5. Billari, F., Aparicio Diaz, B., Fent, T. and Prskawetz A. (2007). The Wedding–Ring. An agentbased marriage model based on social interaction. *Demographic Research*, 17(3): 59–82. <http://www.demographicresearch.org/Volumes/Vol17/3/17-3.pdf>

6. Billari, F. and Prskawetz, A. (2003) (eds.), *AgentBased Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behaviour*. Heidelberg: Physica Verlag.
7. Bühler, Ch. and Fratzczak, E. (2007). Learning from others and receiving support: the impact of personal networks on fertility intentions in Poland. *European Societies*, 9(3): 359–382.
8. Cederman, L.E. (1997). *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*. Princeton, NJ: Princeton University Press.
9. Courgeau, D. (2007). *Multilevel Synthesis: From the Group to the Individual*. Springer Series on Demographic Methods and Population Analysis, vol. 18. Dordrecht: Springer.
10. Courgeau, D. (2012). *Probability and Social Science. Methodological Relationships between the two Approaches*. Methodos Series, vol. 10. Dordrecht: Springer. Forthcoming.
11. Di Piazza, A. and Pearthree, E. (1999). The spread of the 'Lapita people': a demographic simulation. *Journal of Artificial Societies and Social Simulation* vol. 2, no. 3, <http://jasss.soc.surrey.ac.uk/2/3/4.html>
12. 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>
13. Gilbert, N. and Troitzsch, K.G. (2005). *Simulation for the Social Scientist*. 2nd ed. Maidenhead: Open University Press.
14. Hajnal, J. (1955). The prospects for population forecasts. *Journal of the American Statistical Association*, 50(270): 309–322.
15. 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.), op. cit. (pp. 73–96).
16. Keyfitz, N. (1981). The limits of population forecasting. *Population and Development Review*, 7(4): 579–593.
17. Lee, R. D. and Carter, L. R. (1992). Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association*, 87(419): 659–671.
18. Moss, S. and Edmonds, B. (2005). Towards Good Social Science. *Journal of Artificial Societies and Social Simulation* 8(4)13, <http://jasss.soc.surrey.ac.uk/8/4/13.html>
19. Read, D.W. (1998). Kinship based demographic simulation of societal processes. *Journal of Artificial Societies and Social Simulation* vol. 1, no. 1, <http://jasss.soc.surrey.ac.uk/1/1/1.html>
20. Schelling, T. (1978). *Micromotives and Macrobehavior*. Norton, New York, USA.
21. Silverman, E., Bijak, J. and Noble, J. (2011) Feeding the Beast: Can Computational Demographic Models Free Us from the Tyranny of Data? In: T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo and R. Doursat (eds.) *Advances in Artificial Life, ECAL 2011*, MIT Press, Cambridge MA, USA (pp. 747–754).
22. Silverman, E. and Bryden, J. (2007). From artificial societies to new social science theory. In F. Almeida e Costa, L.M. Rocha, E. Costa, I. Harvey and A. Coutinho (eds.) *Advances in Artificial Life, 9th European Conference, ECAL 2007 Proceedings*. BerlinHeidelberg: Springer (pp. 645–654).
23. 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.
24. Wright, G. and Goodwin, P. (2009). Decision making and planning under low levels of predictability: Enhancing the scenario method. *International Journal of Forecasting*, 25(4): 813–825.