SIMULATING THE COST OF SOCIAL CARE IN AN AGEING POPULATION

Eric Silverman, Jason Hilton, Jason Noble and Jakub Bijak
The Care Life Cycle Project
University of Southampton
Southampton SO17 1BJ, United Kingdom
Email: e.silverman@soton.ac.uk

KEYWORDS

Policy Modelling, Social Simulation, Agent-Based Modelling, Social Care, Population Change

ABSTRACT

In this paper we present an agent-based model of the ageing UK population. The goal of this model is to integrate statistical demographic projections of the UK population with an agent-based platform that allows us to examine the interaction between population change and the cost of social care in an ageing population. The model captures the basic processes which affect the demand for and supply of social care, including fertility, mortality, health status, and partnership formation and dissolution. The mortality and fertility rates in this population are drawn from statistical demographic projections until 2050 based on UK population data from 1951 - 2011. Results show that, in general, we expect the cost of social care in the UK to rise significantly as the population continues to age. An in-depth sensitivity analysis performed using Gaussian Process Emulators confirms that the level of care need within the population and the age of retirement have the most profound impact on the projected cost of social care.

INTRODUCTION

As the UK population continues to age, the shift in the age structure of the population puts ever-increasing strain on the country's social care infrastructure. The elderly are the primary consumers of social care services, and dropping birthrates combined with lengthening lifespans mean that while the demand for social care continues to increase, we can expect the supply of social care to decrease as the workforce also ages (Coleman 2002).

The Care Life Cycle Project is pursuing an interdisciplinary approach to this critical societal problem (Brailsford et al. 2012). Given the multiple complex life-course transitions that can affect the provision of social care – decisions in partnership formation, changes in health status, internal migration, and many others – understanding the population dynamics underlying this issue requires an approach that can capture the interactions between these factors. While statistical demographic analyses are critical in order to show us the population change we can expect in the coming decades, we must supplement these figures with a more detailed examination of the processes underlying these dynamics in order to make substantive policy recommendations on the distribution and organisation of social care services.

We propose that combining agent-based approaches with empirically-driven statistical demography allows us to better align social simulations with the 'real world'. In addition, such models can benefit the study of population change by harnessing the flexibility of agent-based approaches to allow for exploration of *scenarios* of population change. This allows for a shift in demographic studies toward explaining and understanding the processes underlying population change, rather than focusing entirely on the prediction and description of empirical trends (Silverman et al. 2012).

Our model moves in this direction by presenting an agent-based platform which captures the complex life-course transitions that influence the demand and supply of social care (Noble et al. 2012). Agents are embedded in a 2D space designed to reflect UK geography, and during the course of the simulation they go through processes of fertility, mortality, partnership formation and dissolution, household formation, and internal migration. Combining this approach with empirical projections until 2050 and vital registration data from 1951-2011 has provided us with a platform that combines the explanatory power of agent-based modelling with the macro-level predictive power of statistical demography (Silverman et al. 2011).

THE MODEL

Basics of the model

Space precludes a complete description of the many parameters present within the simulation. We have attempted to provide a comprehensive overview here; for those who wish

to examine the model more closely, you can find the annotated Python code available at http://users.ecs.soton.ac.uk/jn2/software.php. Code for the current version incorporating our demographic modifications will be made available as well.

In order to facilitate the use of UK demographic data, we based the model's spatial environment on a rough representation of UK geography. Agents occupy houses, which are then grouped into towns which consist of clusters of up to 625 houses. The size of these clusters vary according to local population density which varies across the 8×12 grid which represents the UK.

The agent population has a scaling factor of 1:10,000, given that modelling the entire UK population of 62 million individuals would be prohibitively expensive computationally. This is a greater reduction in scale than is used in microsimulation models of populations, given that the spatially-embedded nature of the agents coupled with their complex behaviours requires substantially more computer power.

The model runs on time steps one year in length. The initial population is generated and distributed randomly in simulation year 1860, then the model runs until 2050, at which point the final figures for social care cost are collected. The model starts early in order to ensure that the population dynamics of the agents can settle prior to the integration of empirical data into the model in 1951.

Agents are able to form and dissolve partnerships with one another. We use the term 'partnership' to refer to any possible relationship that may produce children. Every year agents not currently in partnerships will enter the marriage market, which operates on a national level. Agents are paired with available opposite-sex agents if they meet each other's criteria for a mate. Partnership dissolution is driven by an annual age-specific probability that the male partner will leave. The parameters governing partnership behaviour are not, at present, derived from empirical data, but are instead estimates based on observed rates of partnership formation and dissolution.

Health status and care need

Agents in the simulation begin in a normal health state, in which we assume they require no additional care services. There are age- and sex-specific probabilities which may push the agents into a different care need category; these are checked annually amongst all agents in the simulation. Agents can transition into any of several different levels of care need, shown in Table I. Once agents transition into a care-need state, their health status may only degrade further – they do not recover.

The model also investigates the supply side of social care, by linking the provision of informal care to the availability of agents and their household structure. We assume that any agent will provide care to

TABLE I: The different care need categories, with the number of hours of care required per week

Care need category	Weekly hours of care required
None	0
Low	8
Moderate	16
Substantial	30
Critical	80

any member of their household who requires it, so long as they have time available. Agents have varying amounts of time available with which to provide care, depending on their current status: dependent children can provide five hours per week; adults still living at home can provide 30 hours; and retired people can provide 60 hours. Agents who require care themselves can provide care as well, but only if their own care-need status is Low, and even then only for half the normal number of hours for their status.

The current model does not explicitly represent formal care institutions, such as care homes or similar. Instead, we assume that all available agents (agents in the same household, or children living in the same town) will provide care to the best of their ability, and any remaining care needs are provided by the formal care system. We assume that the state is able to provide this formal care at a cost of £20 per hour. This figure is a rough estimate, to be supplemented by current data when this becomes available. Inflation is not represented in this model, so all care costs are expressed in 2012 UK pounds.

Agent life-course

Newborn agents are classified as dependent children, transitioning to adulthood at the age of 17. All agents enter the workforce at this time and become tax-payers. Adults are further classified into those who have moved out of the family home and those who have chosen to stay with their parents. Agents who reach the age of 65 retire from the workforce and cease paying tax.

The model allows agents to migrate to a new house on the map under several different conditions. When agents form a partnership, there is a 30% chance that they will choose to form a new household with their partner. Agents who take this option have a 30% probability of moving into their partner's home, which may still be the family home of that partner and thus contain other family members. Agents who choose to move elsewhere will either form a new household in the same town or an adjacent town to the location of one partner. When a partnership dissolves, the male agent will move elsewhere on the map, while any dependent children resulting from that partnership will stay with the mother.

Agents can also migrate to new households independently of an partnership upon reaching adulthood. There is an annual age-specific probability
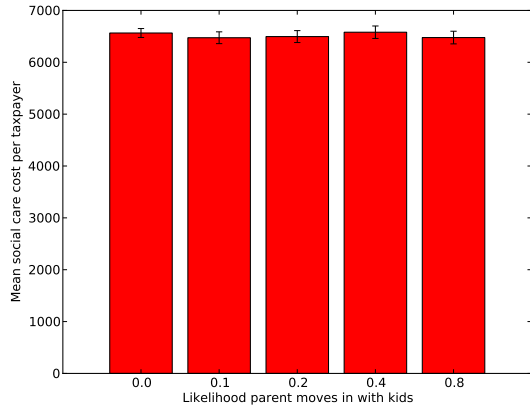
Fig. 1: Results for five different values of the parameter controlling the likelihood of aged parents returning home to live with their children.

Fig. 2: Results for five different values of the parameter specifying the hours of informal care provided by retired individuals to their family members.

that an agent will make this choice. On occasion, single adults or family groups may make an arbitrary move – representing perhaps a change in career or life circumstances that would require a move. Retired agents who live alone also have a small annual probability of moving in with one of their surviving children, scaling inversely with the distance between their current town of residence and the location of the child.

Finally, on rare occasions dependent children agents will have both parents die prior to them reaching adulthood. In these cases we have the agent adopted by a randomly-chosen couple, and the child will move to join their household.

Demographic Projections

The original version of this model used a simple Gompertz-Makeham mortality model which was tuned to give reasonable mortality rates for a modern industrialised nation (Noble et al. 2012). Fertility rates were represented simply by a flat probability of reproduction for any woman of reproductive age in a partnership. In order to increase the realism of this iteration of the model and to tighten its integration with empirical data, we replaced these simplifications with more robust models of mortality and fertility. The shift toward realistic mortality projections ensures that we capture the complexities of the trend toward longer lifespans, and the incorporation of a realistic and detailed fertility model captures the societal move toward later child-bearing and lower birth-rates. These additions were inspired by previous work combining agent-based models with statistical demography (Silverman et al. 2012).

In the new mortality model, we continue to use the approximations used in the previous iteration until the simulation reaches 1951, at which point we switch to age-specific mortality rates drawn from the Human Mortality database 2011 until simulation year 2009. 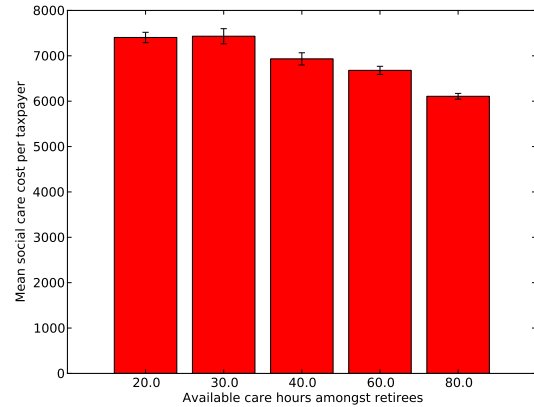Similarly, in 1951 the fertility model switches from a single rate to age-specific mortality rates drawn from the Office of National Statistics data 1998 for England and Wales (for simulation years 1951-1972), then the Eurostat database 2011 for UK women of childbearing age (for simulation years 1973–2009).

These rates were projected forward until 2050 using the forecasting method developed by Lee and Carter 1992. The Lee-Carter 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 Lee and Carter (1992). The forecasts performed through 2050 using this method show life expectancy continually increasing over the period, though the increase slows gradually.

Fertility rates used in this model were also based on empirical data. Age-specific fertility rates from 1973–2009 for UK women of childbearing age were obtained from the Eurostat database (2011), while earlier data for the period 1950–1972 were taken from the Office of National Statistics data for England and Wales (1998).

Once again we used a Lee-Carter model to obtain future fertility rates. In this case two components of the singular value decomposition matrix of fertility rates were used, as two time indices are required to capture the trends in fertility. Projections to 2050 using this method show an initial rise in total fertility rate prior to a convergence at a rate just above replacement fertility. Overall, we see a continuation of the current empirical trend toward later childbearing.
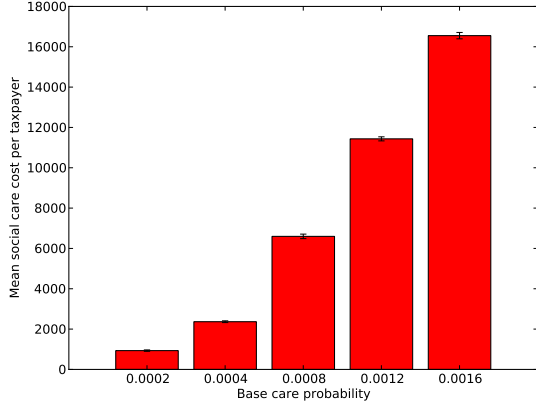
Fig. 3: Results for five different values of the parameter specifying the base probability that an agent transitions to requiring social care.
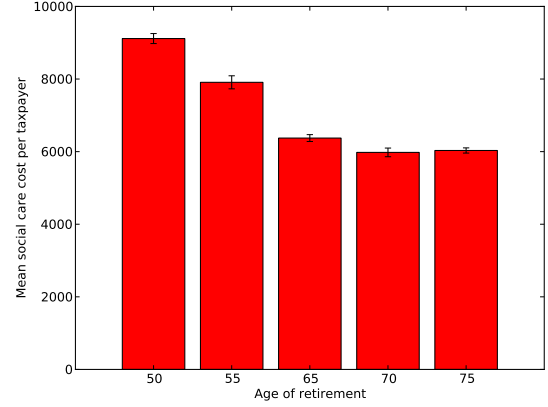

Fig. 4: Results for five different values of the parameter specifying the retirement age for agents within the simulation.

RESULTS

Basic results

Initial investigations of the model were driven by our goal of determining which parameters have the greatest impact upon the cost of social care per taxpayer in 2050. We ran the model under five different variations of four critical parameters for ten runs each, measuring as output the mean social care cost per taxpayer per year at the end of the simulation. In the graphs detailed here, the vertical axis represents care cost, while the horizontal axis displays the mean care cost for ten runs of a given parameter value. The centre bar of each histogram is the default value of that parameter.

The parameters we investigate – the probability of retired parents moving in with children, hours of informal care that can be provided by retirees, base probability for social care need for an individual agent, and retirement age – were chosen because they are areas which policy changes could have some influence (i.e., via changes in formal care provision, institutionalised care provision, tax and welfare changes, etc.). The five variations of each parameter were chosen to give a reasonable range of values from very conservative to very optimistic estimates.

Figure 1 shows the results of a series of runs in which we varied the probability of retired parents returning home to live with their children. We assumed that this may have an impact on the final social care figures, as these retired adults more accessible to informal care from their children. Instead, we found that even significant alterations in this parameter have no perceivable impact on the final figures. This suggests that the share of retired parents moving back in with their children has no significant impact on final care cost figures.

Figure 2 shows the results of varying the number of hours of informal care that can be provided by retired adults. Here we see that high levels of care

availability amongst this group do appear to impact the final social care costs, at least at the higher settings. Retirees can provide the largest amount of informal care hours, and as a group are more likely to live in a household with someone who needs care, so a greater availability of informal care in these circumstances appears to take some burden off the costs to the state.

In Figure 3 we see the results of varying the base probability for an agent to transition into a state of care need. This base probability is then modified by the age and sex of each agent; this allows us to represent the increased need for care amongst elderly males as compared to females, for example. Varying this parameter has a dramatic impact on the final care costs; doubling the default value produces social care costs that balloon to far more than double the results we see at the default level.

Finally, in Figure 4 we examine the impact of changing the retirement age within the simulation. The default age is 65, and here we varied that age between 50 and 75 in increments of 5 years of age. Results show a significant reduction in overall care costs when agents retire later in life; this appears to be due to the substantial benefits of having a larger tax base as agents stay in the workforce longer and continue to pay into the system.

Interestingly, this pattern appears to level off at the age of 70, indicating a point of diminishing returns – and at age 75 the cost of social care actually rises slightly. This occurs despite the lack of any modelling of the health or wellbeing impact that may come from working until very late in life. We suspect that this appears in the model due to the reduced availability of care from adults remaining in the workforce; at a certain point, the increased tax payments from elderly workers will be offset by the reduction in available informal care amongst households with elderly members who are more likely to require care.
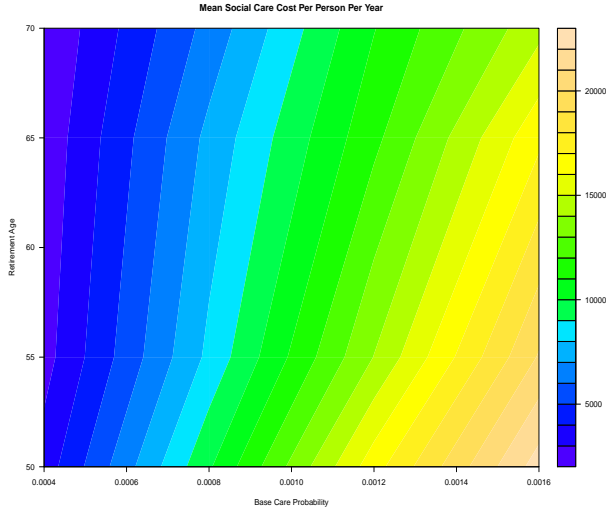
Mean Social Care Cost Per Person Per Year

Fig. 5: Filled contour plot showing the model outputs for varying values of the base care probability and retirement age. The plot shows that a low retirement age combined with high care need probability generates the highest social care costs, producing mean costs per taxpayer well above £20,000 per year.

Input name	Variance (%)
Parents Moving In	0.01
Base Care Prob	88.99
Retired Hours	0.88
Retirement Age	8.06
Parents Moving In.Base Care Prob	0.00
Parents Moving In.Retired Hours	0.01
Parents Moving In.Retirement Age	0.01
Base Care Prob.Retired Hours	0.05
Base Care Prob.Retirement Age	1.88
Retired Hours.Retirement Age	0.11
	Total = 99.9876

Fig. 6: Results of a sensitivity analysis performed using GEM-SA software. Results show that probability of care need and retirement age account for the vast majority of variation in the results. Source: GEM-SA software (own calculations).

These results indicate that retirement age and base care-need probability have the strongest impact on model outputs. Figure 5 provides a concise illustration of the impact of these two parameters. High care-need probabilities combined with low retirement age produces some extremely high values for the mean social care cost per taxpayer per year. In contrast a combination of high retirement age and low base care-need probability produce much more manageable figures, as greater numbers of healthy agents are continuing to work and provide care informally well into later life.

Sensitivity Analysis

The results above indicated that some of these parameters are interacting in interesting and perhaps non-obvious ways, in particular the results for altering retirement age and available informal care hours amongst the retired. As part of the Care Life Cycle Project we have been investigating methods for quantifying uncertainty within statistical and computational models, and this model provides an excellent test case for some of these methods, given the complex interactions occurring between the systems represented here.

Here we took inspiration from another project, Managing Uncertainty in Complex Models (`http://www.mucm.ac.uk/`) and used specialised software to build a statistical emulator of the computational model. These emulators are Gaussian processes analysed in a Bayesian framework, allowing us to produce an effective sensitivity analysis which provides insight into the relative importance of various input parameters in the final output variance (O'Hagan 2006). These methods also provide a means of ac-

counting for uncertainty within the program code, using an additional term called a nugget (Kennedy 2004).

Space precludes a detailed description of the statistical methods underlying Gaussian process emulators, so instead we refer the interested reader to Kennedy and O'Hagan (2001). In brief, these emulators function by assuming that the output variable (in this case, mean social care cost per taxpayer per year) can be decomposed into a constant (mean) term, a series of main effects related to particular input parameters, and a series of interaction effects for all possible combinations of input parameters. The final result is a measure of how much of the total output variance is accounted for by each individual input and each possible combination of inputs.

In our case, we used the four parameters investigated above as our input parameters for the emulator, and our output measure was the mean social care cost per taxpayer per year in simulation year 2050. Running multiple sets of runs at all possible combinations of the values shown above for those parameters produced a training set of some 1,300 results, which were then fed into the emulator. The emulator and the resultant analysis were produced by the GEM-SA (Gaussian Emulation Machine for Sensitivity Analysis) software version 1.1 by Kennedy (2004).

The results in Figure 6 show that the vast majority of the output variance is accounted for by our alterations in the base probability of an agent requiring care (88.99%), followed by the retirement age (8.06%). In contrast, the amount of hours available from retired carers and the probability of retired adults moving home had very little impact (accounting for 0.88% and 0.01% of the output variance, re-
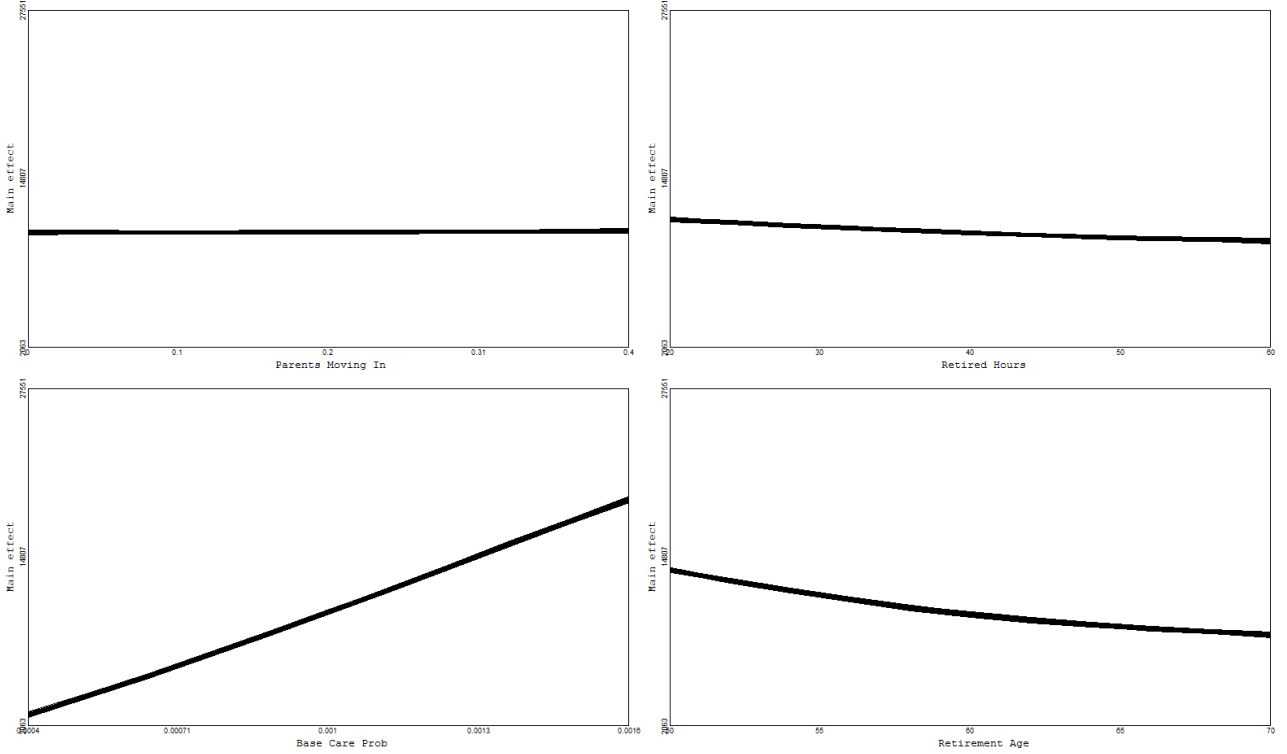
Fig. 7: Results of Gaussian Process Emulator demonstrating the impact of each parameter on final output values. The output value is the mean cost of social care per taxpayer per year at the end of the simulation in year 2050. Source: GEM-SA software (own calculations).

spectively). These results confirm our findings presented in Figure 3 and Figure 4. Figure 7 provides a visual representation of emulator results; each graph shows the main effects of each parameter, with each individual line representing one emulation run.

There was some additional impact from the two-way interactions between these parameters, most noticeably in the interaction between retirement age and the base probability of care need (1.88%). In general however the interaction effects were dwarfed by the impact of the main effects.

CONCLUSIONS

As suggested in our previous work with this modelling framework, we bolstered the realism of our model using significantly more detailed mortality and fertility models derived from statistical demographic data (Noble et al. 2012). The incorporation of Lee-Carter projections (1992) has moved us further toward a useful integration of demographic data and agent-based approaches, and future work will continue to push this forward, as proposed in Silverman et al. (2011).

The results and sensitivity analysis shown here provide a few hints as to the future challenges facing an ageing population. We see dramatic shifts in the results when we alter the retirement age and the probability of agents needing social care. This suggests that two major 'policy levers' which policymakers may examine would be comprehensive pro-

grammes for health amongst the elderly population, and a small increase in the retirement age. The former could significantly alter care need levels by providing preventative care, thus reducing long-term social care costs. This could also allow for the provision of more informal care in this age group, as healthier partners would be better able to provide care – and given that most social care in the UK is provided by family members (Vlachantoni et al. 2011), this could have a significant impact. The latter would allow a significant increase in the tax base amongst older citizens, allowing for more funding for social care costs.

However, our results indicate that the benefits of increasing retirement age level off at around 70 years of age. Intriguingly, this occurs even without any explicit representation of the health impact of working until late in life. Future iterations of the model could shed more light on this effect by allowing for more richness in agent decision-making and incorporating the affects of education and socioeconomic status on health in old age.

More generally, future work can refine these predictions by adding more detail to the social care structures in the model, in particular by adding a more detailed representation of formal care provision. We also will add a more robust migration element which allows for international migration, thus representing the impact of young workers entering the UK. Demographic studies have shown that some level of replacement migration in tandem with policies aimed

at increased labour force participation and raising birth rates might ameliorate the challenges of an ageing society (Bijak et al. 2008).

In summary, our model has provided a useful platform for discussion around the issue of social care provision in an ageing UK society, and has shed some light on the interactions between retirement age, informal social care provision and overall social care cost. The combination of these illustrative scenarios and in-depth sensitivity analyses gives us a strong suite of tools with which to examine the potential impact of policy decisions and economic and social shifts on social care provision in an ageing society.

ACKNOWLEDGMENTS

REFERENCES

Bijak, J., Kupiszewska, D., and Kupiszewska, M. (2008). Replacement migration revisited: Simulations of the effects of selected population and labor market strategies for the aging Europe, 2002–2052. *Population Research and Policy Review*, 27(3):321–342.

Brailsford, S., Silverman, E., Rossiter, S., Bijak, J., Shaw, R., Viana, J., Noble, J., Efstathiou, S., and Vlachantoni, A. (2012). Complex systems modelling for supply and demand in health and social care. In Jain, S., Creasey, R., Himmelspach, J., White, K., and Fu, M., editors, *Proceedings of the 2011 Winter Simulation Conference*. IEEE.

Coleman, D. (2002). Replacement migration, or why everyone is going to have to live in Korea: a fable for our times from the United Nations. *Philosophical Transactions of the Royal Society B*, 357(1420):583–598.

Eurostat (2011). Eurostat statistics database: Domain population and social conditions. http://epp.eurostat.ec.europa.eu. Accessed 27/10/2011.

Human Mortality Database (2011). Human mortality database. http://www.mortality.org/cgi-bin/hmd. Accessed 26/07/2011.

Kennedy, M. (2004). Description of the Gaussian process model used in GEM-SA. Software manual. http://ctcd.group.shef.ac.uk/gem.html. Accessed 23/05/2012.

Kennedy, M. and O'Hagan, T. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society, Series B*, 63(3):425–464.

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

Noble, J., Silverman, E., Bijak, J., Rossiter, S., Evandrou, M., Bullock, S., Vlachantoni, A., and Falkingham, J. (2012). Linked lives: the utility of an agent-based approach to modelling partnership and household formation in the context of social care. In Laroque, C., Himmelspach, J., Pasupathy, R., Rose, O., and Uhrmacher, J., editors, *Proceedings of the 2012 Winter Simulation Conference*. IEEE.

Office for National Statistics (1998). *Birth Statistics, Series FM1 (27)*. Office for National Statistics, London.

O'Hagan, A. (2006). Bayesian analysis of computer code outputs: a tutorial. *Reliability Engineering and System Safety*, 91(10-11):1290–1300.

Silverman, E., Bijak, J., Cao, V., and Hilton, J. (2012). Semi-artificial models of population: Connecting demography with agent-based modelling. In *Proceedings of the 4th World Congress on 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., editors, *Advances in Artificial Life, ECAL 2011*, pages 747–754. MIT Press, Cambridge, MA.

Vlachantoni, A., Shaw, R., Willis, R., Evandrou, M., and Luff, J. (2011). Measuring unmet need for social care amongst older people. *Population Trends*, 145:60–76.

AUTHOR BIOGRAPHIES

ERIC SILVERMAN is a Research Fellow on the Care Life Cycle project at the University of Southampton. He received his PhD from the University of Leeds, and previously worked as a JSPS Postdoctoral Research Fellow at the University of Tokyo. His work on the CLC Project focuses on the application of complexity science approaches to the social science domain. His email address is e.silverman@soton.ac.uk.

JASON HILTON is a Post-Graduate Researcher at the Institute of Complex Systems Simulation and the Division of Social Statistics and Demography at Southampton. He holds a BA in Politics from the University of York, and an MSc in Demography from the University of Southampton. His PhD work focuses on applications of complex systems simulation in demography. His email address is jdh4g10@soton.ac.uk.

JASON NOBLE is a Lecturer in Computer Science at Southampton. He received a DPhil in Cognitive and Computing Sciences from the University of Sussex in 1998. He is the Taught Programme Director for Southampton's Institute for Complex Systems Simulation and has a background spanning AI, philosophy of science, psychology, and statistics. His email address is jn2@ecs.soton.ac.uk.

JAKUB BIJAK is a Lecturer in Demography at Southampton. He holds a PhD from the Warsaw School of Economics, obtained for work on Bayesian migration forecasting. His research interests encompass the applications of quantitative methods in demography, with focus on demographic uncertainty, population forecasting, migration and demography of conflict. His email address is J.Bijak@soton.ac.uk.