Chapter 11
Model-Based Demography in Practice: II

Eric Silverman, Jason Noble, Jason Hilton, and Jakub Bijak

11.1 Introduction

The previous chapter provided some examples of the practice of model-based demography. We were able to develop a simulation which, despite significant simplifications, was able to illustrate some of the core concepts of model-based demography. The model captured the core demographic processes and their role in partnership formation, while remaining free of excessive and expensive data demands. Using Gaussian process emulators allowed us to investigate the impact of key model parameters, making the model more tractable and more useful as a potential policy-making tool.

Now we will examine another simulation which builds upon these foundations. By increasing the complexity of the modelled agents and focusing on a specific policy question, we will illustrate how agent-based modelling combined with statistical demographic projections can create a useful platform for experimentation with social policy. Once again the use of Gaussian process emulators will provide further insight on our results, demonstrating the unexpected interactions that can occur in complex, interlinked social processes.

11.2 Model Motivations

This modelling project came about as part of the EPSRC-funded 'Care Life Cycle' project at the University of Southampton, which ran from 2010 to 2015. This project sought to develop innovative means for predicting the eventual supply and demand of social care in the United Kingdom, using methods drawn from complexity science (Brailsford et al. 2012). The project team consisted of a core group of academics drawn from a variety of disciplines, including demography, gerontology, operations research, and agent-based modelling.

E. Silverman, *Methodological Investigations in Agent-Based Modelling*,
Methodos Series 13, https://doi.org/10.1007/978-3-319-72408-9_11

As with many industrialised nations, the age structure of the population of the United Kingdom has been shifting over recent decades. Birth rates are low and life spans continue to lengthen, resulting in an ever-increasing demand for social care services (Raphael Wittenberg et al. 2051). This trend is becoming increasingly severe as the supply of social care also begins to decrease due to the ageing of the care workforce (Coleman 2002).

Projections of future care demand and supply have for the most part been performed using statistical methods, as in the citations above. However, the Care Life Cycle acknowledged that the provision of social care is inherently a complex process, involving factors present at multiple levels of society. Within families, informal social care decisions can involve aspects of partnership status, socioeconomic status, social networks, and more; similarly, formal care involves issues of the cost of care provision, the availability of carers and migrant workers, and complicated links between public and private providers. A full understanding of these processes the dynamics of the interactions between them requires us to move beyond statistical approaches alone and apply the methods of complex systems science (Silverman et al. 2012).

In this model we address this critical area of public policy using an agent-based model combined with demographic projections. As we saw in the previous chapter, this kind of approach allows us to take advantage of the power of agent-based models and their ability to represent complex processes while keeping our results closely aligned to empirical data and real-world demographic processes. By generating possible scenarios of future social care demand and supply, and analysing the results using uncertainty quantification methods, we can better understand the effect of the social factors and processes underlying this problem and develop an effective platform for testing potential policy responses.

11.3 The 'Linked Lives' Model

This model is an agent-based, spatially-embedded platform in which simulated individuals live out their life-courses in a virtual space roughly modelled on UK geography (Silverman et al. 2013). The version presented here is an extension of a previous model, dubbed the 'Linked Lives' model, which demonstrated the core components of this framework (Noble et al. 2012).

11.3.1 Basic Model Characteristics

The simulation was written from scratch in the programming language Python, chosen both for its ease-of-use even for inexperienced programmers and the availability of convenient libraries for data analysis and plotting (Noble et al. 2012). Python is an interpreted rather than a compiled language, meaning that simulation runs are not optimised for any specific CPU architecture and thus slower; however, the average simulation run can still finish in approximately 1–1.5 min if real-time visualisations are turned on, or 20–30 s if they are turned off.[1] The simulation code is freely available and is provided with the MIT Licence, which is the most permissive open-source licence.[2]

Agents in our model occupy a rough analogue of UK geography divided into a grid which consists of virtual ‘towns’ each containing up to 625 households. The number of households allocated to each town is determined according to the local population density, which is set in the initial conditions of the model to reflect the approximate population density of the UK. Agents represent individuals and can have differing individual characteristics including age, sex, location, work status and health status; the health status aspect of the model will be detailed further below. The simulation operates at roughly a 1:10,000 scaling factor compared to the real UK population, meaning that an agent population representing the UK at the time of writing would consist of approximately 6,400 agents. This scaling was done in order to reduce the computational requirements of the simulation, as the activities of 60 million agents or more would take very large amounts of computer time to calculate.

Agents live out their life-courses in one-year time steps, at the end of which their movements and changes in status are calculated and recorded in detailed log files at both the individual and population levels. In order to allow time for the simulated population to ‘settle’, we begin each run in the year 1860. A simulation run ends in the year 2050, at which point we recorded the projected values for social care cost.

The agents are capable of forming and dissolving partnerships in the model, though note that this simulation does not use the social pressure model of partnership formation outlined in the previous chapter (please consult Noble et al. 2012 for the specifics of this aspect). For the purposes of this simulation we only model partnerships as relationships that can produce offspring. The marriage market for agents is nationwide, and agents will be paired up if they both meet one another’s partnership criteria and if neither already has a partner. Partnership dissolution is a simple process driven by age-specific probabilities of the male agent leaving the partnership.

11.3.2 Health Status Component

In contrast to the Wedding Doughnut, agents in the Linked Lives model are able to transition into varying levels of social care need. Every year agents are checked

[1]These runs were all performed on a 2009-era personal desktop computer with a 2.8 GHz i7 quad-core processor, 12 GB RAM and a 7200 RPM hard drive. We expect runs would be significantly faster on a more modern machine with a solid-state drive and more up-to-date CPU architecture.

[2]The code can be accessed and downloaded here: https://github.com/thorsilver/ABM-for-social-care

Table 11.1 Care need categories

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

Table 11.2 Care provision for each agent category

Agent status	Weekly hours of care available
Dependent children	5
Adult living at home	30
Retired agents	60

against the age- and sex-specific probability that they may transition into a state of care need. Agents who are already in a state of need may worsen in future transitions, but we assume agents do not improve again once they have entered a state of long-term limiting illness. The different levels of care need and the hours of care required for each are in Table 11.1 below.

Unlike the Wedding Doughnut's simplistic model of health status, in Linked Lives the provision of informal social care is modelled explicitly. Each agent is capable of providing care according to their own health status and the amount of free time they have available, which depends on their stage in the life-course and their work status, as seen in Table 11.2 below. Note that agents who are ill themselves *can* provide care for others, but only if they themselves require 'Low' levels of care. In addition they can only provide half the hours of care normally provided by someone at their stage of the life-course and work status. The simulation assumes that agents will provide informal care if they have time available to any member of their household that enters a state of care need; we do not model varied levels of willingness to care, though this is a target for a future expansion of this simulation.

The social care aspect of the model is further simplified by avoiding an explicit representation of formal care, such as care homes or inpatient treatment. The model adds together any unmet care needs of individual agents and 'charges' these to the state at a rate of £20 /h to arrive at the final figure of social care cost in each time step. The intent of the model is not to predict social care costs accurately, but instead to demonstrate the relative impact of different policy solutions on that figure, so in that sense the realism of that figure was not important to the final analyses to be performed in this version (and obtaining the data for these figures would be a complex undertaking all of its own). Note also that the simple economic model present here does not include inflation, so the final costs would actually be somewhat higher than what is presented here.

11.3.3 Agent Behaviours and Demographic Elements

As agents age during the course of a simulation run, they can undergo several different life-course transitions, each of which is recorded and can affect other aspects of the simulation. When agents are born they are classed as dependent children, and are thus able to provide only a small amount of care, as outlined above. Once agents reach the age of 17 they are considered adults and are able to enter the workforce and begin paying tax into the system. During adulthood agents may choose to live at home with their parents or move out on their own. At the default settings, agents cease working and become retired at age 65, but as we shall see later, changing this parameter has some interesting effects.

The first Linked Lives model used a Gompertz-Makeham model of mortality, and fertility rates were a flat probability of reproduction for any partnered female agents of childbearing age (Noble et al. 2012). In this revision we linked these aspects with empirical data by making use of demographic data starting in the year 1951 (the date of the first UK census) (Silverman et al. 2013). As with the Wedding Doughnut, we used age-specific mortality rates from the Human Mortality Database (2011), and fertility rates from the Office for National Statistics (1998) and Eurostat (2011).[3]

We then used the techniques pioneered by Lee and Carter (1992) to calculate projections for these rates until the end of the simulation in 2050. In the case of mortality, these projections show continual increases in lifespan throughout the simulation period, though the rate of increase slows over time. In the case of fertility, we see the rates converge at just above replacement fertility, and a continuing trend toward later childbearing.[4]

The Linked Lives model represents migration in greater detail than the Wedding Doughnut. Agents are able to migrate under several different conditions. When agents form a partnership, they may choose to move into their partner's home, which may still contain members of that partner's family. They may also choose to move elsewhere and start a new household, which can be in the same town as one partner or an adjacent one. Agents may also move independently without a partner; once an agent reaches adulthood there is an age-specific probability that they may make this move in a given year.

Some agents may also choose to move arbitrarily, rather than in response to a change of status; this is intended to be reflective of other individual life choices, such as changing careers or simply desiring a change. Agents who opt to dissolve a partnership will in turn dissolve their household, with the male agent moving to a new random household, while the mother retains custody of any child agents resulting from their partnership. On those rare occasions when both parents die

[3]The mortality rates from HMD run from 1951–2009, while the ONS fertility rates are used from 1951–1972 and Eurostat rates from 1972–2009.

[4]Mortality rate projections used a singular value decomposition matrix of centred mortality rates. Fertility rate projections used two components of the singular value decomposition matrix in order to best capture fertility trends.

before their children reach working age, the children are adopted by a random household and move to join them.

Upon reaching retirement, there is a small chance that agents may choose to move in with one of their children. This was intended to represent the choice of some retirees to move in with their children when they start suffering more health problems. This aspect is investigated further in the Results section.

11.4 Results

Given the Linked Lives model's focus on providing a platform for investigating policy, analysis of simulation results focused on exploring the effects of key parameters. Several parameters were designed to replicate proposed policy solutions, including increasing the retirement age to increase tax receipts, and encouraging aged parents to move back in with their adult children to receive care. Prior to a more in-depth sensitivity analysis, a series of parameter sweeps were run to investigate the impact of these parameters on the cost of social care per capita in simulation year 2050.

11.4.1 Parameter Sweeps

Each of the graphs in this section uses an identical format, with the final per capita social care cost as the vertical axis, and relevant parameter values along the horizontal axis. The displayed values for each parameter setting are the mean final output values for ten runs at that setting.

Figure 11.1 focuses on a parameter which defines the probability in any given year that an older parent may choose to move back in with their adult children. When building the simulation, we had theorised that encouraging older parents to move in with their children would reduce care costs by making it more likely those parents would have access to formal care. As we can see in the graph, however, even when the probability was eight times higher the social care costs were nearly identical. Ultimately it seems that the share of older agents taking this option was not large enough to produce a noticeable reduction in cost.

In Fig. 11.2 we see the impact of changing the hours of care that retired agents are able to provide for loved ones. Here there is a significant impact on cost as the availability of care increases; retired agents are the most likely to live with another agent, generally a spouse, who is also older and more likely to need care. Thus when these agents are able to give more, a relatively large amount of care need is then taken on by this group rather than passed on to the state.

Next we have an unsurprising result in Fig. 11.3, which shows the final social care costs resulting from increasing probabilities of agents transitioning into a state of care need. As this parameter essentially decides the overall health of the agent population, increasing the likelihood of care transitions has a dramatic impact on cost.

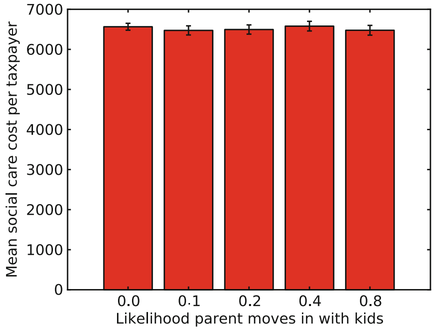

Fig. 11.1 Results of five ten-run series of the model for different values of the parameter controlling the likelihood of parents moving home to live with their children

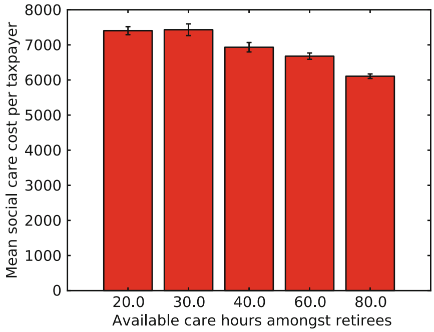

Fig. 11.2 Results of five ten-run series of the model for different values of the parameter specifying the hours of care provided by retired family members

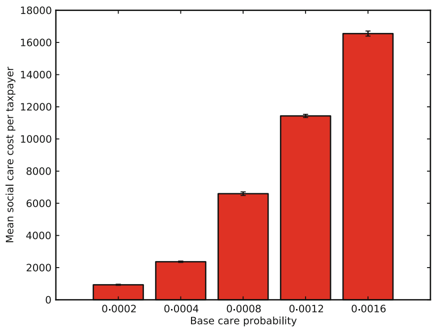

Fig. 11.3 Results of five ten-run series of the model for different values of the parameter defining the base probability of transitioning into a state of care need

Finally, Fig. 11.4 shows the impact of increasing the agent retirement age. Costs increase significantly when the retirement age is reduced, which fits results seen in the real world. However, while increasing the retirement age does reduce cost to a point, the effect levels off beyond age 70. This suggests that keeping agents in the workforce longer does help to some degree, but it also reduces the availability of informal care for older people significantly, as agents who become sick may no longer be able to fulfil their care needs by living with a retired spouse with time to spare.

Interestingly, this result is consistent despite the simulation lacking any means for representing the possible health impact of working later in life. In that respect the model may actually be presenting an optimistic picture in this case. A future

Fig. 11.4 Results of five ten-run series of the model for different values of the parameter specifying the retirement age for working agents

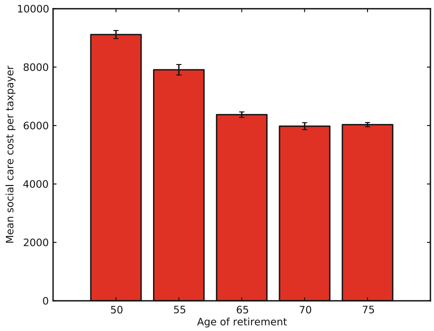

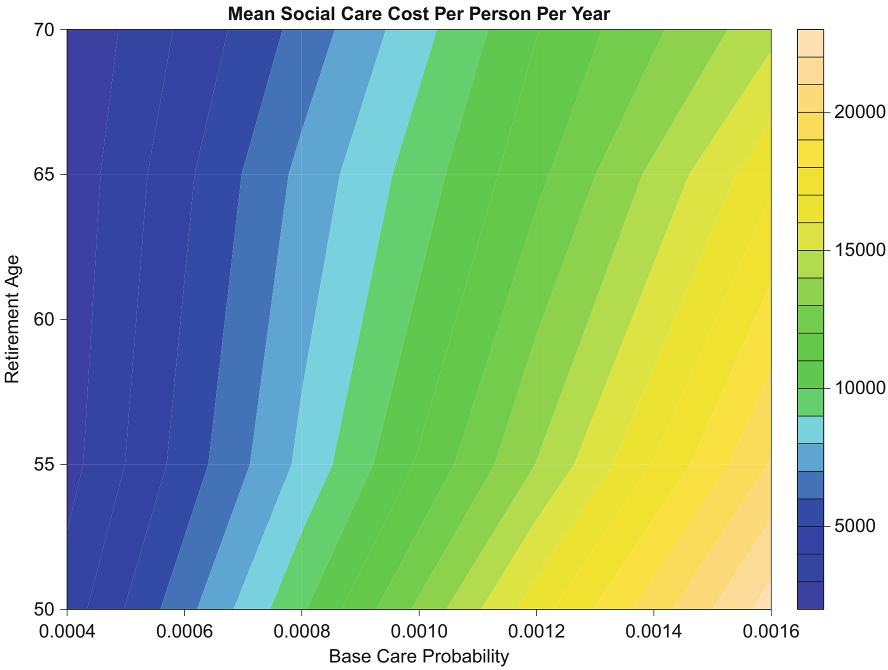

Fig. 11.5 Filled contour plot displaying model outputs for ranges of values for base care probability and retirement age. High probability of care need transitions combined with low retirement ages produces extremely high social care costs

version of the model will examine this aspect more closely and incorporate more detail in this critical part of the life course.

The contour plot in Fig. 11.5 provides a synthesis of sorts, illustrating the impact of retirement age and the base probability of care need transitions on the final social care cost. Scenarios featuring a low retirement age and generally poor population health produce extremely high per capita care cots, reaching well over 20,000 per year. On the other end of the scale, the lowest costs appear in scenarios with high population health and high retirement age; in these scenarios we have a larger proportion of healthy agents able to provide care, as well as lower overall numbers of agents needing care in the first place.

Figure 11.5 serves as a useful illustration of how the investigation of a model's parameter space can help us to explore possible population scenarios. By running the model several hundred times across a range of values for these two parameters, we are able to see clearly how these two values interact. We would not want to use these scenarios as firm predictions, but when dealing with a problem as complex as social care supply and demand these scenario explorations help us to understand the interactions between the policy instruments we have available.

The parameter sweeps above show that even in a simplified model the processes underlying social care supply and demand can interact in unexpected ways. The retirement age investigation in particular illustrates that the model actually produced a surprising insight: that a significant amount of savings are generated for the public purse through the care provided by elderly relatives. These results have since been confirmed by a Carers UK and Age UK report and subsequent follow-up research into what they call the 'invisible army' of older carers who save the UK an estimated 5.9 billion a year on social care costs (Age 2016, 2017). The results from this model pre-date that report by three years, which suggests that well-constructed agent-based models can help us to identify and anticipate significant trends in population health even before relevant data collection has occurred.

11.4.2 Sensitivity Analysis with Emulators

The complex, interacting processes at play in the social care system are represented in a relatively simplistic way in this simulation, but as we have seen, the results still demonstrate some interesting dynamics. In order to better understand the relationships between these four key model parameters, a Gaussian process emulator (O'Hagan 2006) was used in a similar way to the Wedding Doughnut model.

In this case, the final output of interest is the final social care cost per capita in simulation year 2050. The four input parameters discussed above were used in the emulator, and a range of parameter values were chosen to provide a reasonable section of the parameter space to investigate. Once again we followed the example of Kennedy (2004) and used the additional 'nugget' term to account for additional uncertainty generated by the program code itself.

In order to provide a substantial dataset for the emulator, multiple runs of the simulation were performed at all possible combinations of the chosen parameter values. This resulted in approximately 1,300 simulation runs, the outputs of which were condensed into a pair of text files and fed into the GEM-SA (Gaussian Emulation Machine for Sensitivity Analysis) software (Kennedy 2004).

GEM-SA produces two primary forms of output, displayed here in Figs. 11.6 and 11.7. In Fig. 11.6 we have a summary of the percentage of the final output variance accounted for by the parameter named in the left-hand column; the top four rows show the variance resulting from the four main parameters individually,

Fig. 11.6 Visualisation of sensitivity analysis performed with Gaussian process emulator. Results were generated from 41,000 runs of the emulator (Source: GEM-SA software (own calculations))

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

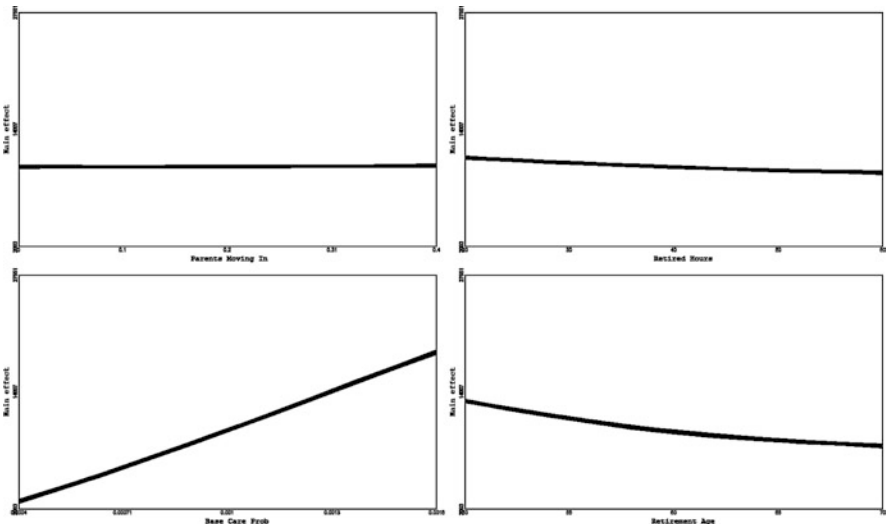

Fig. 11.7 Results of Gaussian Process Emulator demonstrating the impact of each parameter on final output values. The output value is the mean cost of social care per taxpayer per year at the end of the simulation in year 2050 (Source: GEM-SA software (own calculations))

while the six rows beneath record the impact of the interactions between pairs of those parameters. The 'Base Care Prob', or base probability of transitioning to a state of care need, clearly has the most substantial impact by far, accounting for 88.9%; a generally healthy population produces generally lower costs, and vice versa. Conversely, the probability of parents moving in with their children (listed as 'Parents Moving In') has a very small impact on care costs, accounting for only 0.01% of output variance alone and a further 0.02% in interaction with other factors.

However, once again retirement age plays a larger role than anticipated. Some 8.06% of the output variance is accounted for by the effect of the retirement age parameter, with a further 2% in interaction effects. This is significantly higher than the impact of the number of hours of care provided by retired agents (‘Retired Hours’), at 0.88% and 0.17% respectively. Figure 11.7 illustrates the impact of the four parameters on final output variance, and once again this visualisation confirms the surprising impact of retirement age on per capita social care cost.

11.5 The Power of Scenario-Based Approaches

The results of this simulation demonstrate the power of agent-based modelling when applied to a policy-relevant research question. By augmenting an abstracted spatial model of social care (Noble et al. 2012) with demographic data, we are able to investigate the impact of possible policy changes on social care costs. While the model lacks significant details regarding some of the specific processes and social factors underlying the social care system in the United Kingdom, the dynamics were captured well enough to produce results that highlighted the role of older carers in the system, an aspect of the problem that has only begun receiving attention very recently (Carers and Age 2015; Age 2016).

By modelling both the supply and demand of care in the system, the model also allows us to study the impact of policy changes and anticipate possible unintended consequences of these changes. In this case, we can see that while shifting the retirement age upward does prompt significant reductions in care costs, this is not a bottomless well of additional tax revenue. Once we raise the retirement age beyond a critical point, the savings evaporate as older carers are then less available to provide informal care for their loved ones, the cost of which is then passed on to the state. Future iterations of the model could produce more specific recommendations in this respect by explicitly modelling the health impact of both longer careers and significant caring responsibilities amongst older people. As noted by Age UK, older carers are more likely than non-carers to report their health as ‘not good’, and are more likely to report being anxious or depressed (Carers and Age 2015), so modelling the pressures on this critical group can help us to better target support systems designed to help carers cope.

Speaking more broadly, this model takes us another step down the road toward a model-based demography. By enhancing a relatively simple agent-based model with detailed mortality and fertility projections derived from real data, we are able to construct a simulation that integrates the strengths of statistical demography with the flexibility of simulation, all without relying on excessive amounts of empirical data. Further refinements to the approach will enable us to connect our models more closely to empirical data, enhancing the predictive capacity of the models, while also maintaining the critically important balance between realism, generality, precision, and tractability.

The model can certainly be criticised for not providing point predictions of specific cost estimates, or fine levels of detail for different parts of the UK. However, the theoretical backstory of the model is designed to justify its use as an exploration of possible scenarios related to social care cost in a simulated ageing population. The lack of predictive capacity in this case is a conscious design decision, as we are seeking an understanding of the relative impact of certain factors on care costs and an exploration of possible policy impacts, not an economic forecast based on highly-detailed datasets. So while we may not be able to offer these results to the Department of Health as unambiguous recommendations of specific actions to take, even this simple model can produce some useful insights regarding the limited benefits of increasing retirement age in response to the pressures of an ageing population with growing social care needs.

Future expansions of this model could focus on aspects such as international migration, a particularly important element in this context as demographic research has shown that policies aimed at replacement migration and increasing birth rates can help reduce the impact of ageing trends (Bijak et al. 2008). Additional details on the pressures facing older carers, including the impact of working in later life, would help to generate a more detailed portrait of that critical part of the informal care structures evident in UK society. Initial versions of the model were intended to include some effects derived from socioeconomic class, which is well known to influence health outcomes for older people (Majer et al. 2011), but this was removed for the sake of simplicity; future versions could include this aspect to allow the model to reflect the complex impact of health inequalities in the UK.

In summary, the model provides an early-stage example of the value of scenario exploration in the context of policy-relevant demographic models. The use of Gaussian process emulators further enhances the model's power by allowing us to pick apart our results and identify the relative impact of key parameters on the final output of interest. This kind of scenario-based approach allows us to look beyond the one-generation time horizon and anticipate the possible long-term impacts of population and policy changes. Future versions of this model will continue to push this approach forward by modelling more relevant social factors.

In the next chapter we will take stock of our progress in model-based demography thus far, and discuss some key challenges that face the field in the future.

Acknowledgements This work was supported by the UK's Engineering and Physical Sciences Research Council, grant EP/H021698/1 'The Care Life Cycle'. We thank Marc Kennedy and Anthony O'Hagan for providing the powerful GEM-SA software to the public at no cost.

References

Age UK. (2016). Invisible army of oldest carers saving state billions. http://www.ageuk.org.uk/latest-news/archive/invisible-army-of-oldest-carers-saving-state-billions/. [Online; posted 18 May 2016].

Age UK. (2017). *Briefing: Health and care of older people in England 2017*, London.

Bijak, J., Kupiszewska, D., & 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., & Vlachantoni, A. (2012). Complex systems modelling for supply and demand in health and social care. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, & M. Fu (Eds.), *Proceedings of the 2011 winter simulation conference*. IEEE.

Carers UK, & Age UK. (2015). *Caring into later life: The pressures on older carers*, London.

Coleman, D. A. (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 Oct 2011.

Human Mortality Database. (2011). Human mortality database. http://www.mortality.org/cgi-bin/hmd. Accessed 26 July 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 May 2012.

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

Majer, I. M., Nusselder, W. J., Mackenbach, J. P., & Kunst, A. E. (2011). Socioeconomic inequalities in life and health expectancies around official retirement age in 10 Western-European countries. *Journal of Epidemiology and Community Health, 65*, 972–979.

Noble, J., Silverman, E., Bijak, J., Rossiter, S., Evandrou, M., Bullock, S., Vlachantoni, A., & 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 C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, & J. M. Uhrmacher (Eds.), *Proceedings of the 2012 winter simulation conference*. IEEE.

Office for National Statistics. (1998). *Birth statistics, series FM1 (27)*. London: Office for National Statistics.

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., & 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., Hilton, J., Noble, J., & Bijak, J. (2013). Simulating the cost of social care in an ageing society. In W. Rekdalsbakken, R. T. Bye, & H. Zhang (Eds.), *Proceedings of the 27th European conference on modelling and simulation* (pp. 689–695). Dudweiler: Digitaldruck Pirrot.

Wittenberg, R., Comas-Herrera, A., Pickard, L., & Hancock, R. (2004). Future demand for long-term care in the UK. A summary of projections of long-term care finance for older people to 2051. *Health and Ageing, 11*, 2–4.