Different Ways of Modelling Phone Adoption

Lynne Hamill

Centre for Research in Social Simulation, Department of Sociology, University of Surrey, Guildford, Surrey, GU2 7XH, United Kingdom
lynne.hamill@surrey.ac.uk

Abstract. Systems dynamics and agent-based models are used here to examine the spread of fixed line phones in Britain over 120 years. Both models approximately fit the data and it is shown that in this case the two approaches can be used to complement each other. The SD model is simpler and produces a better fit while its deterministic nature facilitates sensitivity analysis. The agent-based model provides greater explanatory power, which can in turn be used to fine tune the systems dynamics model. Together, they can be combined to tell a plausible story about the adoption of telephones.

Keywords. Systems Dynamics. Agent-based modelling. Technology adoption.

1. Introduction

This paper follows Scholl's recommendation [1] that system dynamics (SD) and agent-based (AB) models be compared on identical topics. The relative strengths of the two models are explored by looking at the adoption of fixed line phones in Britain from their introduction in 1880 to the end of the twentieth century. The aim of the paper is to explore methodological issues rather than to provide a robust explanation of the observed adoption pattern, which will be the subject of later work.

Section 2 provides the historical background to the adoption pattern that the models are to be tested against. The theory is set out in Section 3. Sections 4 and 5 describe the SD and AB models respectively. The models are compared in Section 6. Section 7 shows ways in which the two approaches can be used to complement each other and Section 8 concludes.

2. History

Britain's first public telephone exchange opened in London in August 1879 to serve eight subscribers. By the end of the year the number of subscribers in London totalled 200 and exchanges had been opened in seven other cities. The first phone directory was issued in 1880 and “contained details of over 250 subscribers” plus details of 16 provincial exchanges [2]. By 1882, there was one phone for every 3,000 people in London: by 1890, the ratio was up to one in about 800 but it did not reach one per 100 until 1905 [3]. However, these numbers included business as well as private subscribers and it is likely that the adoption rate for households was very much lower, even in London, and was lower still outside London. However, the data series for households only starts in 1964, by which time 21.6% of British households had fixed line phones. This percentage peaked at 95% in 1999, after which it started to drop as mobiles were substituted [4].

Thus phones took some 80 years to spread from virtually no households to around 20% and almost 120 years to reach 95%. Perry [3] suggests that this slow take-off was due to price and poor regulation of the nascent phone industry. Also, there was limited geographical coverage: London's first trunk line was not opened until 1884 and it was not possible for Londoners to call “the Midland and Northern Counties” until 1890 even though the first line to Paris opened in 1891 [2]. Rural areas were not well covered and even by 1913, one third of all telephones were in London [3]. Another complicating factor was that public phones were available from 1886 [2]: this means that it was not necessary to have a phone at home in order to make a call, but person-to-person calls using a public system were, of course, difficult to manage.

3. Theory

The models presented in this paper focus on two important factors that underlie the adoption of phones: the network effect and affordability.

The network effect. In his 1969 seminal paper Bass [5] argued that except for first adopters, take-up of new “generic classes of products” (as opposed to new models of older products) is related to the number of previous buyers. If the new product happens to be a link to a communication network, this effect is particularly important. Metcalfe’s Law states that “the value of a communications network is proportional to the square of the number of its users” and “the law is said to be true of any type of communications network” [6]. Essentially “the idea is that a network is more valuable the more people you can call” [6]. Fischer [7] noted that when phones were introduced in the US they were used “to widen and deepen existing social patterns rather than to alter them”. Valente [8] argued that there are two processes at work: one reliant on the “entire social system” and the other on “an individual’s personal network”. The first is a matter of following changes in society in general, such as opinion leaders who are not personally known, while the other implies that an important determinant of phone adoption is whether your family and friends, that is, those in your ‘personal network’, already have phones.

Affordability. Initially, only the better off could afford phones. In the US "the more affluent households were the earliest subscribers" [7]. It was the same in Britain. Phones were expensive: in 1901 "When you could employ a maid for £20 per year, having unlimited phone use for £17 per year did not seem to be a bargain" [2].

4. A Systems Dynamics Model

The SD model is based on Verhulst's logistic equation [9]:

$$dp/dt = rp(1 - p)$$

where p is the proportion of adopters, t is time and r is a parameter controlling the speed of adoption. This equation means that the rate of adoption is determined by the existing proportion of adopters. Thus this SD model can be said to model the network effect implicitly: the greater proportion of households that have phones, the more likely any given non-adopting household will adopt.

The adoption curve is therefore determined by two factors, a start point and the growth rate, r. As explained above, we only know that the telephone service started in 1879 and took some 80 years to reach about a fifth of households. Although by 1882 there was one phone per 3,000 people in London, this ratio is clearly too high for households because it includes business phones and takes no account of the very few phones outside London. In the absence of better data, it was therefore assumed that one household in 10,000 were adopters in 1880. On this basis, a Verhulst equation with r set at 10.3% provides the best fit (using OLS) for the whole period: but for the period from 1960 to the mid-1970s setting r at 10% gives a better fit, while for the later years 10.5% is better. Also, the SD model gives no reason for the take-off of adoption after 80 years other than a simple network effect i.e. the more people who have a phone the more attractive it is for others to have one too. Nor does it give any indication why the growth rate should be around 10%. (The results are shown in Fig. 1 later in the paper alongside the results of the AB model to facilitate comparison.)

5. An Agent-Based Model

In this AB model 10,000 agents are spread randomly across a toroidal grid of just over 99,000 cells. There are two types of agents: 'Blues', who represent the affluent early subscribers and are all located in one quadrant; and 'Greens', who represent the rest of the population and are spread randomly throughout the whole 'world'. These assumptions are designed to reflect both the affordability and the concentration of adopters both socially and geographically noted above.

To be consistent with the SD model, one of the Blues is designated the first adopter so that one in 10,000 had a phone. In each time period the adopting Blues 'persuade' another Blue within a fixed radius, representing their network of family and friends, to adopt. It was not presumed necessary for all the Blues to adopt before the Greens start adopting because, depending on the distribution of Blues, it may be that not all

Blues would adopt. Instead, it was assumed that once the adoption rate among the Blues stops rising, adoption spreads to the Greens on the same basis.

The percentage of Blues in the population will, by definition, be small and information on personal networks suggests that the number of people contacted frequently is measured in tens (see e.g. [10]). The model was therefore run with the percentage of Blues at 2.5%, 5%, 7.5%, 10% and 15% and the personal network radius increasing in increments of 1 from 5 to 10. (Note this is not the size of the personal network but the radius defining the network.) Each of these 30 combinations was run 10 times, giving a total of 300 runs, because the output of each run varies due to the different random distribution of agents across the 'world'. The average of each group of 10 runs was taken to facilitate comparisons. Table 1 indicates that only six combinations produced a curve that approximated the actual series by producing an adoption rate of 21% by 1960.

Table 1. Predicted adoption rate by 1960 using the AB model (%.)

Personal network radius	**Percent Blues** 2.5	5	7.5	10	15
5	14	9	7	6	6
6	**18**	10	9	6	10
7	**19**	10	7	8	46
8	**20**	11	8	45	77
9	**21**	8	60	82	97
10	**19**	**21**	95	100	100

Closer examination of the results from the six combinations shown in bold in Table 1, including taking OLS, revealed that only when it was assumed that 5% were Blues and the agents had a personal network radius of 10 did the curve approximately match the observed pattern. Typically, this radius implied a personal network of about 30 agents. To confirm this result a further 30 runs were undertaken using these same assumptions. Fig. 2 shows the results of all 40 runs (grey), with the average (a solid black line) and one standard deviation (dashed lines) alongside actual take-up (black squares).

6. Comparing the Models

Figs. 1 and 2 show the results of the models. So which method provides the best model for phone adoption in Britain? Parunak et al [11] argued that AB modelling "is most appropriate for domains characterized by a high degree of localization and distribution and dominated by discrete decisions" and that the choice between the two approaches should be made on a case-by-case basis. But on what basis should that choice be made?

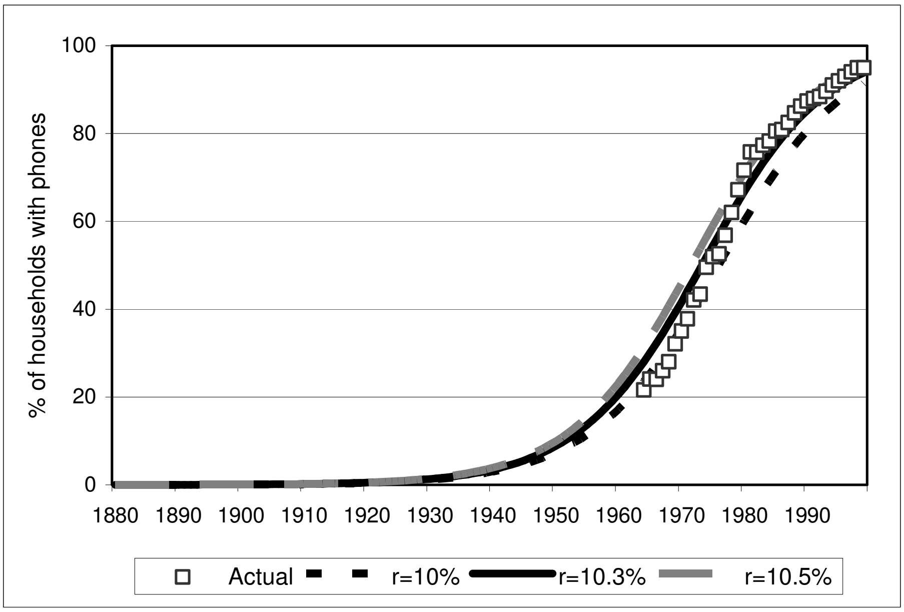

Fig. 1. Household adoption of phones: predictions of the SD model compared to actual.

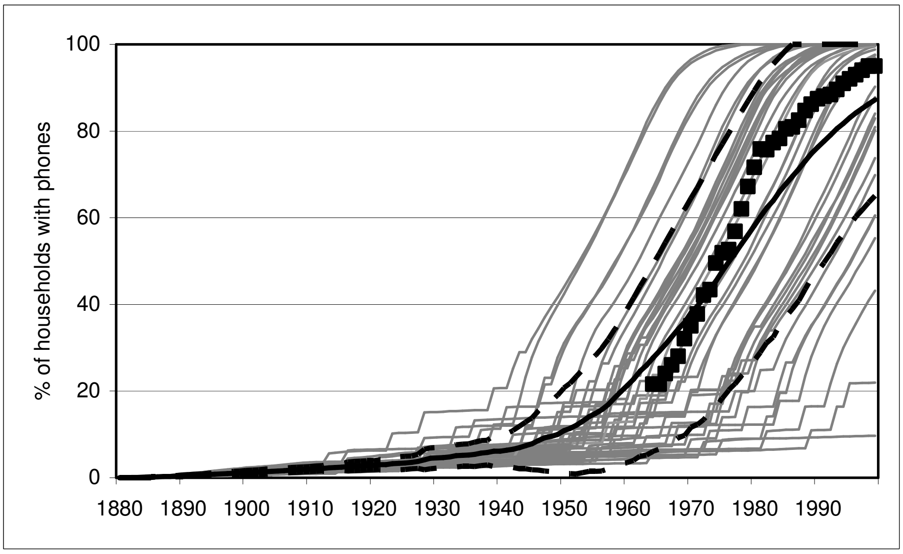

Fig. 2. Household adoption of phones: 40 predictions (grey), the average (solid black line) and one standard deviation (dashed lines) of the AB model assuming 5% Blues and a personal network radius of 10 compared to actual (black squares).

> "A standard modelling principle is that the level and complexity of a model should be chosen so that it answers the questions and embodies the theoretical elements we are interested in, but is otherwise as simple as possible" [12].

How can that principle be applied? I suggest three basic criteria can be used: goodness-of-fit, fitness-for-purpose and simplicity. These criteria are now applied to the two models described above: the SD model with *r* set at 10.3% and the AB model with Blues set at 5% and the personal network radius set at 10.

Goodness-of-fit. The SD model provides a unique result for each set of parameters. However, the results for the AB model vary between runs because the distribution of the agents across the 'world' varies and for this reason, the average is used to measure goodness-of-fit. Fig. 3 shows that in both models it takes some 80 years for adoption to reach about 20%. But the SD model replicates the rise in the following 40 years more accurately. Overall, the SD has a much lower OLS, as illustrated in Fig. 3. Thus the SD model provides the best fit although fine-tuning the parameter values with the AB model might improve its fit.

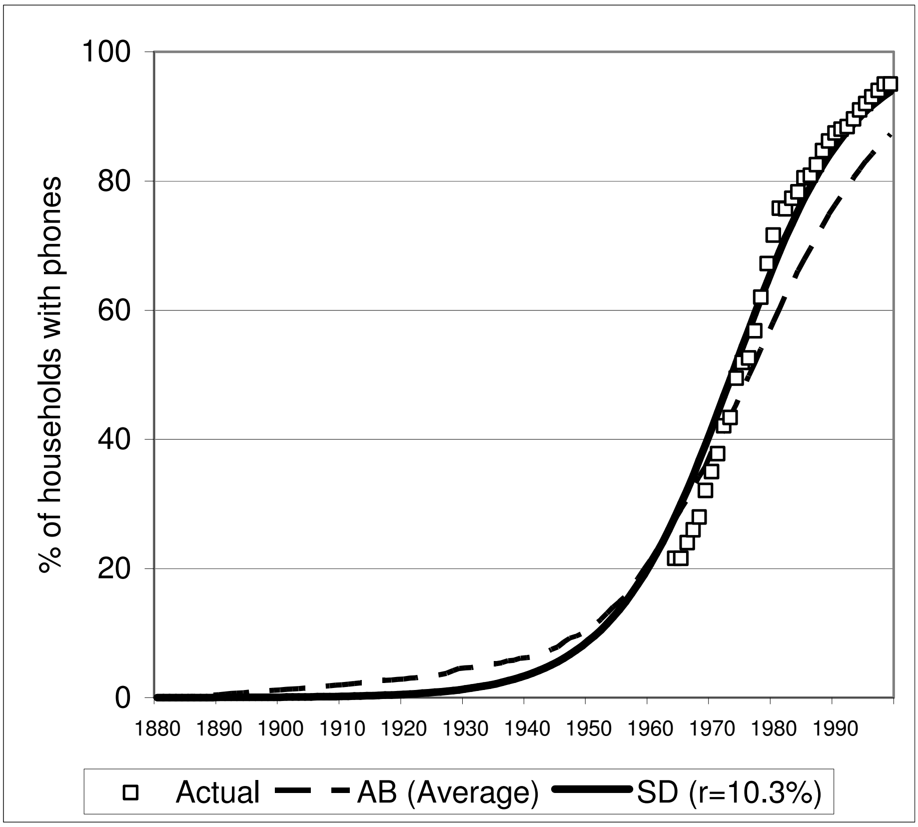

Fig. 3. Comparison the SD and AB models with the actual data.

Fitness-for-purpose. As Forrester pointed out long ago [quoted in 13] the validity of a model should be judged by its suitability for a particular purpose. Many models are built for forecasting or planning. However, a model that predicts well may not add much to our understanding. Quite accurate short term economic forecasts can be made by looking at the time trend of the variable in question and extrapolating it a little way ahead, with no understanding of what is underlying the forecast changes. Indeed, Coleman [14] observed that:

> "macroeconomic predictions based on leading indicators having known statistical association with subsequent system performance may give better predictions than will economic models based on interactions among parts of the system".

In this case the SD model simply reflects the network effect. It provides no insight to understanding the underlying processes whereas the AB model suggests an explanation, telling "a story" that is consistent with the literature.

Simplicity or Occam's razor: *entia non sunt multiplicanda praeter necessitatem* or "entities are not to be multiplied beyond necessity" [15]. The SD model is undoubtedly the simpler to describe and simpler to program. It is also runs faster.

To sum up: in this example, SD scores well on goodness-of-fit and simplicity but low on explanatory power. It is therefore not surprising to see that this is the approach chosen by UK phone supplier BT to model phone uptake (see [16]). The AB model scores lower on fit and simplicity but much higher on explanatory power: phone adoption spread through a geographically and socially close group of affluent early adopters before reaching a wider population. While the fit of the AB model may be improved by adding more parameters, too much fine-tuning of this kind could result in reduction of explanatory power by making interpretation difficult. (Some work was also done on cellular automata but it is not reported here as measured against these criteria, the method was no more than a 'poor man's' AB model: it was poorer fit, offered less explanatory value but was not significantly simpler.)

The models were implemented using NetLogo [17] and can be found at:
www.hamill.co.uk/misc/essa07.zip.

7. Combining the Models

The SD versus AB debate seems often to be presented as an either/or choice, top-down versus bottom-up, macro versus micro. Yet as Fishwick [18] noted, models of different types can be combined to answer different types of question about a given process. Möhring and Troitzsch [19] took an SD model and started to break it down, moving it towards an AB model. How this might work in the case of UK phone adoption?

The SD model assumed everyone was identical whereas in the AB model, there were two types of individuals, the affluent early adopters (Blues) and the rest (Greens), and each individual had their own unique personal network. While the SD model cannot readily deal with 10,000 households, it can deal with two types of people: early adopters and the rest. Now the AB model suggested 5% of the population were affluent early adopters and it so happens that assuming the growth rate for these early adopters was 9¼% and then 13% for the others, a better OLS fit to the data can be obtained using the SD model than simply using a constant rate throughout (see Fig. 4). Thus information from the AB model has been used to improve the SD model.

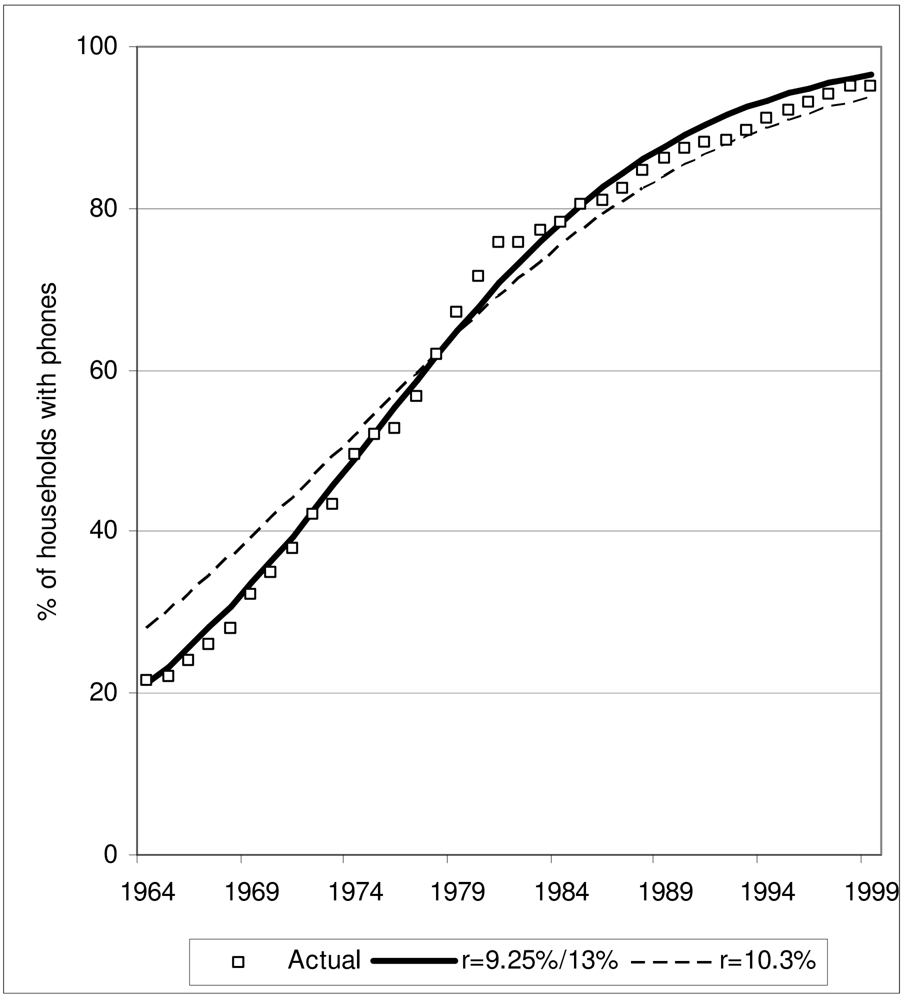

Fig. 4. Household adoption of phones: predictions of the improved SD model compared to the simple SD model and actual.

In turn, the SD model may be used to improve the AB model. The fact that this SD model is deterministic facilitates sensitivity analysis. For example, given the adoption rate of 9¼% for the early adopters and 13% for the rest, varying the percentage of affluent early adopters from 1% to 15% makes a noticeable difference to the adoption rate between 1950 and 1990 but little difference in the early years or as the take-up approaches saturation (see Fig.5). Put another way: the assumed percentage of early adopters, by definition a small group, did not matter much for 70 out of the 120 years being studied!

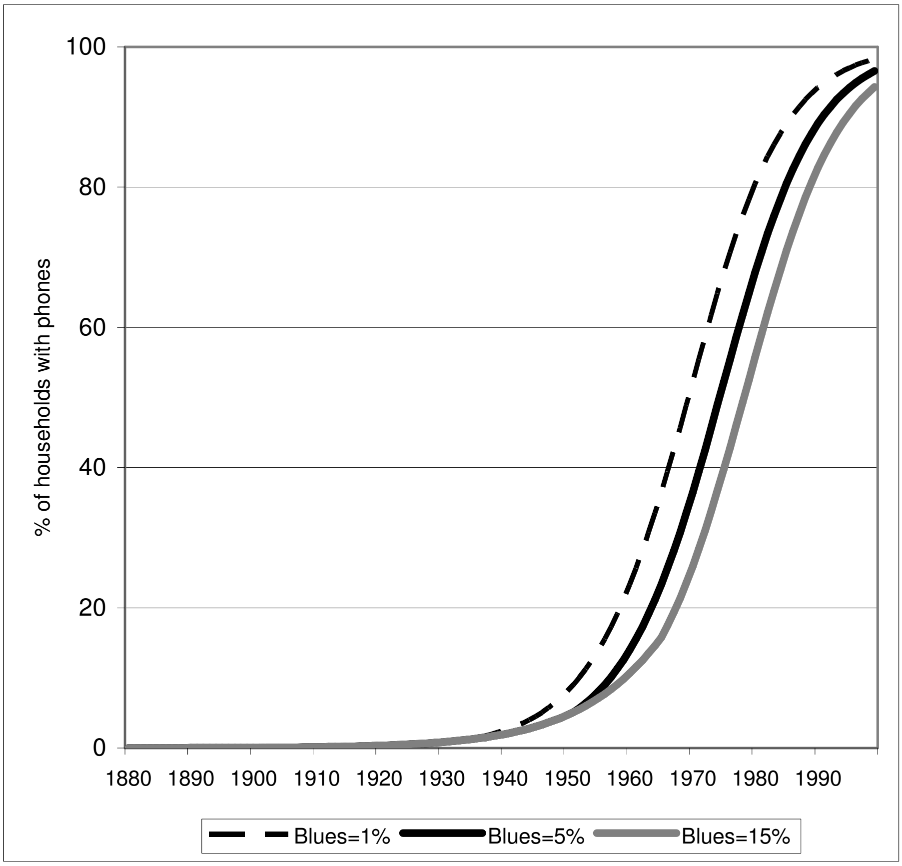

Fig. 5. Results of varying the percentage of early adopters using the SD model.

Because each run of the SD model takes seconds rather than minutes as for the AB model and because using the AB model more runs are needed due to the stochastic processes, it is possible to do more experimenting, more sensitivity analysis with the SD model. For example, the SD model can be used to demonstrate that if 5% of the population were affluent early adopters then unless their growth rate is at least 7% it is not possible for overall take-up to exceed 90% in 120 years (see Fig. 6). In other words, a slow initial take-up can have very long-term consequences! These 121 runs reported in Fig. 6 took under a minute with the SD model: doing the equivalent with the AB model, which would have required perhaps a thousand runs, would have taken many hours, even days.

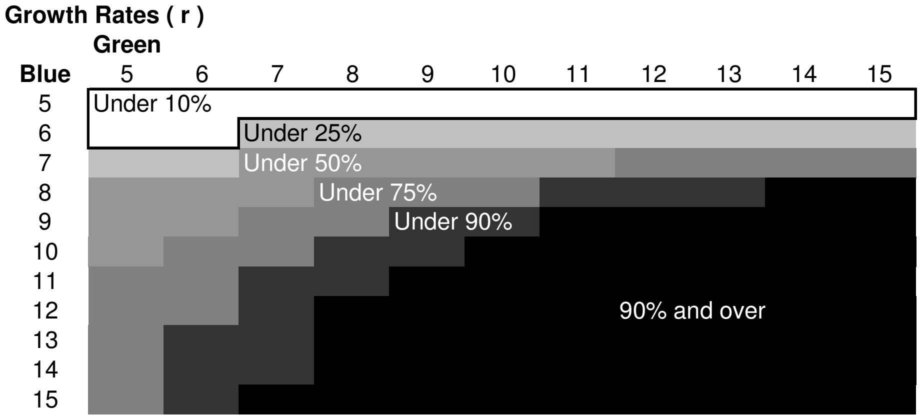

Fig. 6. Results of varying the two growth rates using the SD model.

I suggest that the greater insight provided by the AB model has been used to improve the SD model and that the SD model has been used to test the sensitivity of the assumptions in a way that would have been impractical with an AB model because of its heavy computational demand and its stochastic nature.

8. Conclusion

The main aim of the paper was not to explore in detail the adoption of fixed line phones in Britain. Nevertheless combining the history, theory and the results of the two models an interesting story emerges. The slow adoption of phones for the first 80 years followed by a fast rise to saturation can be explained as follows. About 1 in 20 geographically and socially close affluent households were early adopters. These households had average personal networks of about 30 and they persuaded one member of their network to adopt each year. Once adoption stopped spreading among this group, it started to spread in the same way to the rest of the population. This is consistent with a Verhulst growth rate of about 9¼%, which rises to 13% for the later adopters. With these growth rates, for the first 70 years the outcome is not sensitive to the percentage of early adopters assumed, given that by definition it is a small group. However, if the growth rate for the early adopter group is ‘too low’ – less than 7% on these assumptions – the product will never ‘take-off’. Further work is needed.

More importantly, this paper has shown that SD and AB models can, in some cases, produce similar results and that they appear to complement each other, each having its own strengths and weaknesses. So I suggest that rather than choose between SD and AB models, in some cases, both can usefully be used. Further work will explore the use of SD models to assist formulating and verifying AB models.

Acknowledgements
This work was supported by Microsoft Research through its European PhD Scholarship Programme. Thanks, too, to Nigel Gilbert for advice and encouragement.

References

1. Scholl, H.J.: Agent-based and System Dynamics Modeling: A Call for Cross Study and Joint Research. Pr. of 34th Hawaii International Conference on System Sciences. IEEE. (2001)
2. BT Archives: Events in Telecommunications History: (2007) http://www.btplc.com/Thegroup/BTsHistory/Eventsintelecommunicationshistory/Eventsintelecommunicationshistory.htm Accessed 6 June 2007.
3. Perry, C.R.: The British Experience In de Sola Pool (ed), The Social Impact of the Telephone. MIT Press, Cambridge, Mass.(1977)
4. 1964-93: CSO: Family Spending. HMSO London (1994)
   1994-97: ONS: Family Spending 1997-98. TSO. London (1998)
   1998-99: ONS Family Spending 2005-06. Palgrave Macmillan. Basingstoke (2007)
5. Bass, F.M.: A New Product Growth for Model Consumer Durables. Management Science, Vol 15. No 5, Theory Series. Pp 215-227 (1969)
6. Briscoe, B., Odlyzko, A. & Tilly, B.: Metcalfe's Law is Wrong. (2006) IEEE Spectrum (July 2006) 26-31
7. Fischer, C.S.: America Calling: A Social History of the Telephone to 1940. University of California Press. Berkeley. (1992) 261-3
8. Valente, T.W.: Social network thresholds in the diffusion of innovations. Social Networks, (18) (1996) 69-89
9. Strogatz, S.: Nonlinear Dynamics and Chaos. Westview, Cambridge MA (1994) 22-23
10. Boase, J., Horrigan, J.B., Wellman, B. & Rainie, L.: The Strength of Internet Ties. Pew Internet & American Life Project. Washington (2006)
11. Van Dyke Parunak, H. et al.: "Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide" in "Multi-Agent Systems and Agent-Based Simulation": Book Series Lecture Notes in Computer Science, Springer, Vol. 1534, 1998, 10-25.
12. Doran, J. & Palmer, M.: The EOS Project. In Gilbert, N. & Conte, R. (eds) Artificial Societies. UCL Press. London. (1995) 103-125
13. Qudrat-Ullah, H.: Structural Validation of Systems Dynamics and Agent0Based Simulation Models. Proceedings 19th European Conference on Modelling and Simulation (2005)
14. Coleman, J.S.: Foundations of Social Theory. Harvard University Press, Cambridge, Mass. (1994) 4
15. Oxford Dictionary of Philosophy. OUP. (1994)
16. Lyons, M.H., Burton, F., Egan, B., Lynch, T. & Skelton, S.: Dynamic modeling of present and future service demand. Proceedings of the IEEE, 85(10), (1997) 1544-1555
17. Wilensky, U.: NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. (1999)
18. Fishwick, P. A.: Simulation model design and execution: building digital worlds. Prentice Hall Englewood Cliffs, N.J (1995) 285-332
19. Möhring, M. & Troitzsch K.G.: Lake Anderson Revisited by Agents. Journal of Artificial Societies and Social Simulation 4.3. http://www.soc.surrey.ac.uk/JASSS/4/3/1.html (2001)