4 Mix, Chain and Replicate—Methodologies for Agent–Based Modeling of Social Systems[1]

David Hales

INTRODUCTION

The modeling of social processes and systems using agent-based models (ABM) is increasingly seen as a valid tool over a wide range of disciplines including economics, sociology and anthropology (Gilbert and Troitzsch 2005; Halpin 1999). Although not considered a central tool or method within any one discipline ABM has attracted loyal followers in each area, bringing together researchers from many disciplines with different methodological backgrounds and approaches. This rich mixture of approaches and backgrounds is the primordial soup from which great and original work can evolve but it can also lead to misunderstanding, failure to communicate and, perhaps worst of all, the constant re-emergence of stale and entrenched debates that are very well represented in other areas of social science.

ABM is a technique in which models are composed of a number of subunits called "agents" that represent subentities of a social system. Agents may represent individuals, groups, firms or other entities. In computational simulation work the agents are software abstractions. Agents are represented as algorithms (rule-based decision processes) and data (local agent memory). Agents inhabit a shared environment in which they interact with each other. The scenario being modeled dictates the nature of the agents, their interactions and the environment.

Work in ABM is methodologically permissive. There is no single ABM method or methodology.[2] ABM is a technique or technology rather than a methodology or a discipline. It is important to understand this since it explains why no single methodology would be appropriate for all ABM work.

Methodology is rarely discussed explicitly and in detail in ABM papers because it is assumed that the nature of the investigation and framing of the research questions and the ABM itself should be sufficient for the reader to understand why the particular approach is being employed. This, in general, is the case with good ABM work. However it can be confusing for those new to ABM looking for methodological clarity and can also can lead to confusion between experienced researchers who have used ABM but only from a different methodological tradition.

We identify a number of approaches that can be combined in different ways to reflect many of the methodologies found in the literature. We argue that such approaches can be used to allow ABM researchers to build on each other's findings and models. We believe that finding ways of building on, testing, extending and reapplying findings is a necessary condition for approaching the level of rigor required to support what might be termed a "science" of social systems. Additionally, by linking models from different disciplines and traditions increased communication is possible between different researchers. Communication occurs through the ABM models. The models themselves can become a kind of *lingua franca*.

This chapter is structured in the following way: first we present some quotes from formative researchers in the field concerning methodology, specifically discussing the fact that ABM applied to sociological phenomena incorporate both deduction and induction in interesting and new ways. We then present a "mix and match" approach based loosely on a Popperian (Popper 1968) approach to the analysis of ABM. We then present the idea of "chains of models" and how they relate to ABM. Following this we briefly discuss ABM replication and finally put the pieces together and conclude with some observations on progress in ABM methodology over the last 10 years.

COMBINING DEDUCTION AND INDUCTION

It has been noted by several foundational social simulation researchers that ABM social simulation does not fit neatly into either deductive or inductive methodologies. Consider the following comments:

> Simulation is a third way of doing science. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does not prove theorems . . . induction can be used to find patterns in data, and deduction can be used to find consequences of assumptions, simulation modelling can be used as an aid to intuition. (Axelrod 1997)

> Clearly, agent-based social science does not seem to be either deductive or inductive in the usual senses. But then what is it? We think generative is an appropriate term . . . We consider a given macrostructure to be "explained" by a given micro—specification . . . (Epstein and Axtell 1996)

> We can therefore hope to develop an abstract theory of multiple agent systems and then to transfer its insights to human social systems, without *a priori* commitment to existing particular social theory. (Doran 1998)

> Our stress . . . is on a new experimental methodology consisting of observing theoretical models performing on some testbed. Such a new

> methodology could be defined as "exploratory simulation". (Gilbert and Conte 1995)

In the following section we incorporate these observations into a "mix and match" method combining various components that are found in ABM modeling work, producing many possible kinds of method applicable to ABM depending on the nature of the research questions that are being addressed.

MIX-AND-MATCH METHODOLOGIES

In order to group different ABM approaches into a set of indefinable methods we have imported some Popperian terminology (Popper 1968). Of course here we apply these terms to an artificial deductive system (ABM models) rather than the real world. To be more precise we examine the ABM as an entity "in the world" which can be empirically examined by applying a kind of Popperian approach. We do not claim the approaches we present are exhaustive and we do not wish our tone to be prescriptive. Rather these sketches should be seen as ways to clarify and classify methods already in use in ABM work.

The methods employed in ABM work can be broken down into a collection of reasonably distinct components. These are: a of assumptions (A) that are used to specify the agents and their environment, a set of runs (R) comprising execution of a computer program which embodies A, a set of measurements or observations (O) of the runs, a set of explanations (E) which attempt to link A and O in some meaningful way and a set of hypotheses (H) linked to E based on A and O.

A, R and O are formalized since A is represented by a computer program, R some set of executions of the program and O some specified measures of R. However, E and H may or may not be formalized. They are often given in a mixture of natural language using qualitative concepts and statistical or mathematical relationships. In either case the explanation aims to illuminate the dynamic processes in R with reference to A and O and possibly via the identification of some emergent properties.[3]

Connecting the aforementioned components in different ways reveals several methods of inquiry, some of which are now detailed.

Perhaps the simplest method is the presentation of an *existence proof*. An existence proof does not require E or H at all. Here A is shown to be *sufficient* to produce some O (see Figure 4.1). Much ABM work follows this method, at least in publication presentation, because it is concise and easy to understand. Some assumptions are given and shown to be sufficient to produce some outcome. The evidence is presented based on observations usually shown as charts of individual runs and distributions over multiple runs. In general this kind of method benefits from minimal assumptions (simple agents) and a qualitative easily identifiable outcome.

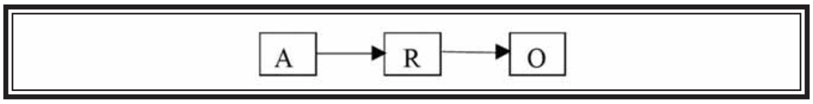

Figure 4.1 The form of an existence proof. The assumptions (*A*) are coded into an ABM, and a set of simulation runs (*R*) support some observations (*O*).

Behavior modeling (or reverse engineering) again does not require *E* or *H*. Here some existing process (*R*') is observed (*O*') and compared (possibly visually/qualitatively) against *O* and, based on divergence, *A* is revised. This process is continued until a satisfactory level of correspondence is observed (see Figure 4.2).

Theory testing involves the translation/abstraction of some existing theory concerning real social processes *T* into *E*, *A* and *H* and then the testing of *H* against *O* in order to either support or refute *H* and by implication *T* (see Figure 4.3). An early example of this was presented by Doran *et al.* (1994) in which a theory of Upper Palaeolithic change was tested.

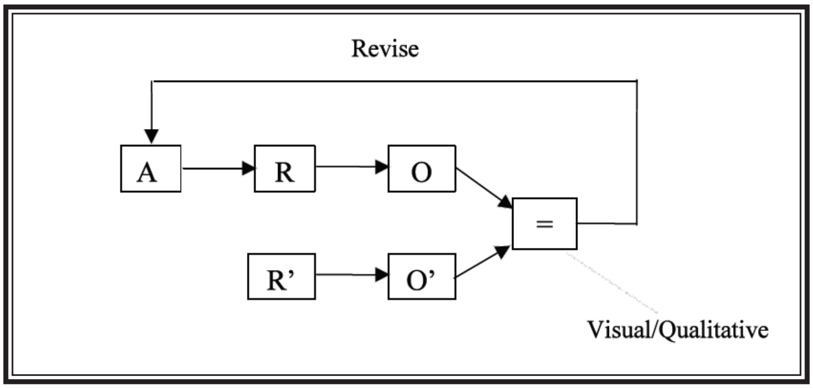

Figure 4.2 Behavior modeling. Observations are compared to some existing process (*R*') producing observations (*O*'). Assumptions are revised to align behavior.

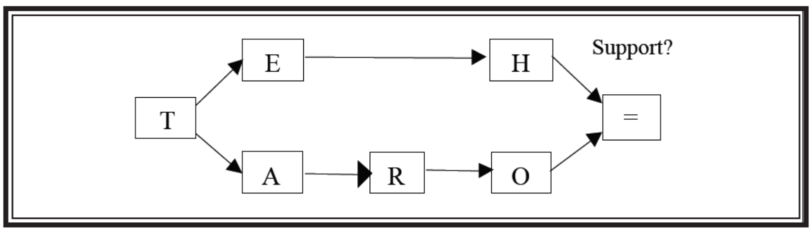

Figure 4.3 Theory testing. An existing theory (*T*) is used to specify assumptions (*A*) and an explanation (*E*) which explains how *A* leads to *O* in the model. From *E* hypotheses are derived (*H*) which predict what *O* should be. These can be tested against *O*.

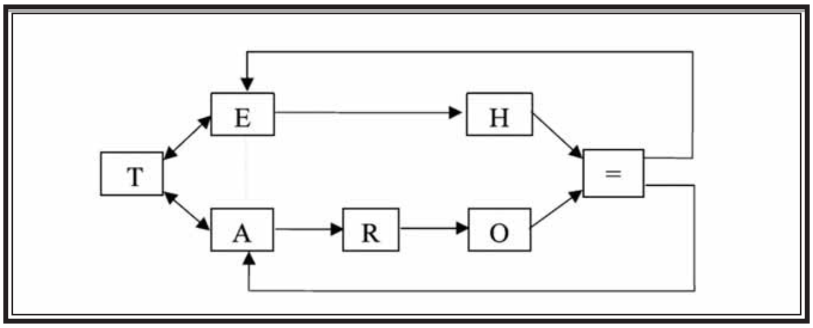

Figure 4.4 Theory testing. By revising the explanation (*E*) and assumptions (*A*) based on finding agreement of hypotheses (*H*) with observations (*O*) new theory can potentially be created.

Theory building involves the abstraction from *T* into *E, A* and *H,* comparison between *O* and *H* and then possible revision of *E* and/or *A*. Given that a state is reached in which *E, H* and *O* correspond, *E* and *H* can then possibly be "de-abstracted" into *T* producing a theory testable against real social processes (see Figure 4.4).

Explanation finding involves iterative refinement of *E* based on comparison of *H* with *O* w ithout changing *A* (see Figure 4.5). This means we fix the assumptions; this might be necessary when the research question involves relatively fixed assumptions which produce *O* of interest but it is not known how this happens—i.e., some emergent property that the ABM modeler, although able to produce, does not understand how *R* produces it. This might be termed "trying to find out what is going on in an ABM by repeatedly applying new hunches and then trying to refute them". Actually this method most closely reflects the spirit of Popper since we cannot change the assumptions (*A*) and we are looking for explanations through a kind of informed trial-and-error process. It is generally the case in this mode that refutation is the easiest course of action by which to test *E*. One can look for some observation that will refute *H*.

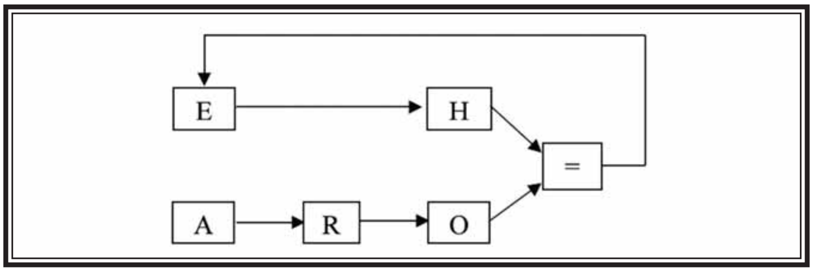

Figure 4.5 Explanation finding. Revise the explanation (*E*) until the derived hypotheses (*H*) match the observations (*O*).

Many of these modes combine deduction and induction often in an iterative way. Such investigation has been termed “ceduction” which is short for “computer experimental induction/deduction” (Hales 1998). The inductive process here is viewed as iterative observation and a revision of *E* and *A*. The *O* is produced deductively (computationally) from *A*, but the revision of *E* and *A* is an inductive process based on observation guided by *H*.

It should be made clear why we have attached strong caveats to our use of the term “Popperian” approach. Although we can use the mix-and match-methods to refute hypotheses (*H*) we can also “change the rules of the universe” by changing *A* to “unrefute” some *H*. This should not be seen as “cheating” but (as we have labeled previously) a kind of theory building or behavior modeling. This is a constructive enterprise in which we ask the question: What assumptions are sufficient to produce certain kinds of observable behavior from the ABM? However it should also be noted that this does not mean “anything goes” because *A* will be constrained by the specific research questions being addressed.

CHAINS OF MODELS

It is often desirable to import certain properties from existing models into new models. For example, a highly abstract model of an artificial society which self-organizes high levels of cooperation between egotistical agents might help to explain a specific target social phenomena if it can be incorporated into a more elaborated and specialized model (by supplementing and possibly changing some assumptions).

But since many of the properties of ABM models result from complex and emergent processes it is rarely easy to identify which elements of the set of assumptions (*A*) are necessary, sufficient or contingent. Hence importing properties from existing models into new models is not a matter of simply selecting known assumptions and combining them with new assumptions.

One way to achieve the import process is to construct chains of models in which the assumptions are varied gradually in each successive model until a sufficient level of detail or abstraction is obtained. The links between models in the chain represent the preservation of some desirable property between models.

Essentially what is happening during an iterative chaining process is that theory, in the form of algorithms evidencing some phenomena of interest, is being carried over into a new scenario or context. This is particularly useful when models are to be moved across disciplinary boundaries. For example, a biologically orientated evolutionary model might display properties that can be used to capture a social process by changing some assumptions or *vice-versa*.

Chains can also be constructed *post hoc,* rather than as part of a goal-orientated process. That is, existing models produced for different reasons and at different levels of detail or application may be found to be chainable if a common link can be found between them—i.e., if they can be shown to share a given property and subset of assumptions that support it. This has been termed model "alignment" or "docking" (Axtell *et al.* 1996) or more generally "model-to-model" analysis (Hales *et al.* 1998). This approach of finding common phenomena and mechanisms operating in different models constructed in different disciplines offers the possibility of finding general and unified underlying processes expressible at different levels. Essentially by linking models in this way one attempts to link or unify theories embodied in the models.

A model chain may terminate when it reaches a target system (real social system) in which it is empirically validated via comparison of the target with the terminal model. We do not discuss in detail how this may be done here but we refer interested readers to the "cross-validation" work of Moss and Edmonds (2005). Essentially cross-validation involves grounding both the assumptions, specifically the micro-behavior of the agents, and the observations of system macro-behavior in real social systems.

Another way that ABM may interact with the real social word is through a construction process that incorporates the stakeholders themselves (the agents being modeled) in the model construction process. This is termed "participator modeling". Again we do not discuss this here as it is covered in detail elsewhere (for a good overview see Ramanath and Gilbert 2004).

More recently ABM social models have been applied to finding engineering solutions, through chains of models, in distributed self-organizing software systems such as agent-based computing (Brueckner *et al.* 2006) and more recently peer-to-peer systems (Hales and Arteconi 2006). In this approach chains terminate when they have reached a level of elaboration required to produce an actual deployable implementation. In this sense validation becomes demonstrating that the software system performs the required functions.

Figure 4.6 shows an example of a chain linking several ABM moving from an abstract social model (TagWorld—Hales 2000) towards two peer-to-peer (P2P) applications: Broadcast (Arteconi and Hales 2006) and CacheWorld (Hales and Marcozzi 2007). The more abstract models are to the left, the more specific to the right. Although both Broadcast and CacheWorld have a common lineage in TagWorld and NetWorld (Hales 2005) they differ considerably as they are modifications of the intermediate models SLACER (Hales 2006) and SkillWorld (Hales 2006). For each model a brief description plus the scenario used are given in the figure. However, these details are not important; rather this is given as an example of model chaining in action. Note that the more abstract models use the prisoner's dilemma game as a test for the emergence of cooperation, and the more applied models emerge cooperation in specific P2P application domains.

It should be noted that chains can run in either direction; for example recent work has taken P2P applications and chained back to new kinds of

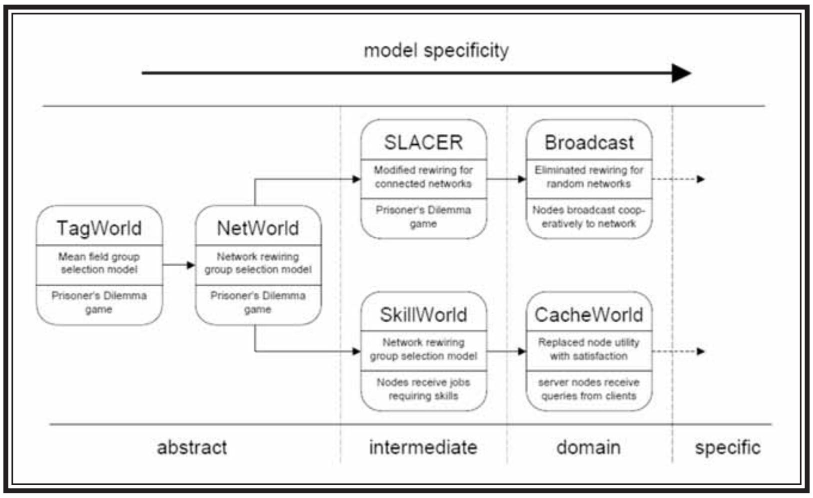

Figure 4.6 Example of a model chain terminating in peer-to-peer application domain models.

social theory models (Mollona and Marcozzi 2009)—these are examples of so-called "peer production" models (Benkler 2006).

REPLICATION

We have argued elsewhere that since ABM are generally not analytically tractable (i.e., we need to use the empirical approaches described earlier) confidence in results can be obtained only via replication of results by independent researchers (Edmonds and Hales 2003). In exactly the same way that empirical findings in scientific areas such as physics need to be replicated to be trusted so do ABM results.

Good replications should ideally work only from the assumptions (*A*) given in the original work. This means ignoring extraneous details such as the specific computational environment, computer languages and tools used since these should not affect the results obtained. Indeed a good replication should start from scratch using different languages and computational abstractions if possible. Essentially the ABM should be recoded based on the assumptions presented in the original work. These assumptions follow a kind of high-level specification, and by replicating the ABM from the specification two critical questions are answered:

- Is the clarity and level of detail of the presented assumptions (*A*) sufficient for an ABM programmer to construct, from scratch, a working model?
- If an ABM can be produced does it replicate the main results and observations (*O*) presented in the original work?

Interestingly, experience has shown that the first question is rarely answerable in the positive and often requires direct communication between the original researchers and the replicators. This should not surprise because original published work needs to follow the space and style constraints of academic publications. In general ABM work is presented as concisely as possible to communicate the general result rather than to give an exhaustive and unambiguous software specification. This can cause serious problems if original authors of work cannot be contacted.

Several ABM researchers have proposed that published work should be supplemented with appendices containing additional detail in the form of a reasonably standardized pseudocode algorithm or flowchart describing the ABM simulation in addition to the original source code (Edmonds 2004; Edmonds and Bryson 2004). However this practice is not widespread at present.

The second question concerning actual reproduction of results is rarely a simple matter of looking for an exact match between observations (*O*) of runs (*R*) in both models. This is because ABM work often involves many runs that produce alternative histories due to stochastic processes (randomness) built into the model. Often then, the issue becomes one of statistical matching of results and/or qualitative matches (i.e., the same emergent phenomena was observed).

In fact the issue of randomness (or more specifically pseudorandomness) pervades ABM. By replicating in other environments different pseudorandom generator algorithms are applied. Experience indicates that it is rare that pseudorandom bias can seriously affect outcomes but it is a possibility.

It has also been noted that rounding errors due to real number representations in digital computers can also lead to seriously misleading results (Polhill *et al.* 2005). Unfortunately, replicated models will often have the same forms of rounding errors since this is a processor or operating system issue rather than an ABM implementation issue. It has been suggested that "interval arithmetic" implementations could be used eliminate this potential source of error, however currently this is very rarely done (but see Polhill and Izquierdo 2005).

Replication can be viewed as a simple and short model chain (as discussed previously). The chain contains two models and the phenomena of interest (to be preserved) are the entire set of observations (*O*) from the preceding model.

It is often claimed that, although desirable, there are few academic incentives to replicate. As we have discussed earlier it is not an easy task and, the argument goes, a positive or negative result does not necessarily lead to quality publications. Reviewers will ask—so what? If you can't reproduce the results perhaps your model is wrong or has a bug,[4] and if you can replicate then what have we learned that is new? However, recently this appears to be changing as ABM become more widely cited and understood (see Will and Hegselmann 2008; Galan and Izquierdo 2005).

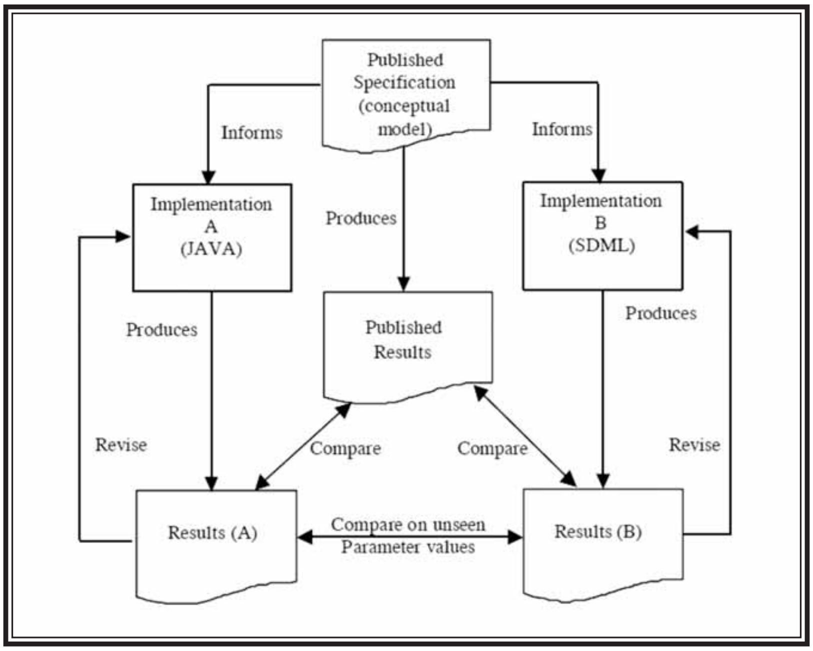

Figure 4.7 Diagram outlining a replication process in which two independent replications were made of a previously published model. The process allowed for a detailed examination of the claimed results of the original model (for details see Edmonds and Hales 2003).

Another incentive for replication comes from using the chaining method discussed previously. If a researcher wishes to apply, say, an abstract model to some more specific domain then the initial work should be to replicate the abstract model in an extensible form before modifying and specializing it. Hence in this way the replication work is a by-product of the chaining process rather than the main focus of the work.

PUTTING IT ALL TOGETHER

If we put together the methods of mix and match, chaining and replication we can think of ABM work as a kind of expanding network of linked models. Nodes represent particular model instantiations (generally reported in some publication); links represent relationships between models (chains and or replications). We can visualize such a network such that nodes on the periphery are more specific and applied and those nearer the core are more abstract and general. That is, nodes at the edge of the network terminate

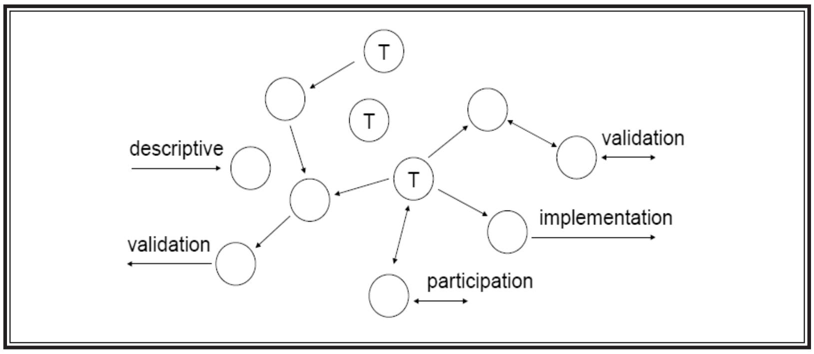

Figure 4.8 A network visualization of models and how they relate. Nodes are ABM models and links represent chain relationships between models. The nodes at the periphery may be seen as linking to empirical social realities through various methods such as empirical validation, engineering implementations and participatory or descriptive processes. Nodes in the center (here marked with a *T*) represent abstract or theoretical models.

where they relate directly to either real-world empirical results (based on a real target system) or, from the engineering perspective, represent instantiation of deployed working software systems.

CONCLUSION

In this chapter we have outlined three broad methods of working with ABM: mix-and-match methodologies, model chains and model replication. We have proposed approaching ABM empirically. We argue that ABM researchers should view their models as aspects of the physical world that can be investigated experimentally like other physical sciences. If analytically tractable and useful models of social behavior can be produced then we do not need to take the ABM route. But it seems evident that ABM researchers should not believe that because they use computer models (based on automatic logical deductions of a computer program) their results are any sounder than those in the empirical sciences. This is experimental science with all the concomitant caveats, pitfalls, opportunities and possibilities.

With this in mind what we have presented in this chapter is a loose summary of a set of methods and approaches that, although diverse, can integrate ABM work from diverse disciplines and with diverse goals. Again looking to the physical sciences we see that it is possible to integrate both highly abstract theory, often based on intuition or mathematical beauty, with empirical experiment and applications. We believe careful use of ABM in social modeling can potentially achieve this through focusing on linking models in chains, replicating important results and using a rigorous empirical methodology towards the ABMs themselves.

Over the last 10 years or so we have observed ABM maturing in a promising direction. We increasingly see physicists working with ABM applying a physics perspective. We see new replications of important models. Also we are seeing work explicitly linking models through model-to-model analysis (Hales *et al.* 2003; Rouchier *et al.* 2008) and cross-fertilization between social models and the engineering of distributed computer systems because the requirements for such systems become ever more social, complex and self—organizing (Di Marzo Serugendo 2007). Recently we have witnessed an explosion of empirical work based on the new and massive data sets available from Internet applications and mobile phone records, and other electronic sources, allowing for levels of detailed social analysis never before possible (Palla *et al.* 2007). This offers potential for large-scale validation of ABM.

We welcome these developments and look forward to the next decade of ABM research.

ACKNOWLEDGMENTS

Many of the ideas and thinking about methodology in this chapter were heavily influenced by Bruce Edmonds and Scott Moss from the Centre for Policy Modelling in Manchester. During my time working there (over four years ago now) extensive methodological discussions were ongoing and I benefited greatly from these. The ideas of replication and chaining are very much directly influenced by Bruce's ideas, work and approach. The earlier "mix and match" ideas were directly influenced by my PhD supervisor (and ABM pioneer), Jim Doran from Essex University (almost a decade ago now). All errors, vagueness and unconvincing arguments are of course my fault.

NOTES

1. This work was partially supported by the Future and Emerging Technologies program FP7-COSI-ICT of the European Commission through project QLectives (Grant no. 231200).
2. We use the words "method" and "methodology" synonymously in this chapter.
3. We do not define or discuss the nature of "emergence" in detail here. The term is used in different ways by different authors. For our purposes it can be considered to mean some observable property that emerges from the runs of a model that is not intuitively expected (or easily reducible) to the assumptions that comprise the rules coded into the agents.
4. One way to address this is to perform a further independent replication to give three models. One can then use a majority vote to determine which model appears to be misbehaving. If all three models disagree we can at least be sure that the specification is too vague to be used for meaningful replication.

BIBLIOGRAPHY

Arteconi, S. and Hales, D. (2006) 'Broadcasting at the critical threshold', *Technical Report UBLCS-2006–22,* University of Bologna, Dept. of Computer Science.

Axelrod, R. (1997) 'Advancing the art of simulation in the social sciences', in R. Conte and R. Hegselmann (eds) *Simulating Social Phenomena—LNEMS 456,* Berlin: Springer.

Axtell, R., Axelrod, R., Epstein, J. and Cohen, M.D. (1996) 'Aligning simulation models: a case study and results', *Computational and Mathematical OrganizationTtheory,* 1: 123–41.

Benkler, Y. (2006) *The Wealth of Networks: how social production transforms markets and freedom,* New Haven, CT: Yale University Press.

Brueckner, S., Di Marzo Serugendo, G., Hales, D. and Zambonelli, F. (eds) (2006) 'Engineering self-organising systems', Proceedings of the 3rd Workshop on Engineering Self–Organising Applications (EOSA'05), *Lecture Notes in Artificial Intelligence,* 3910, Berlin: Springer.

Di Marzo Serugendo, G., Martin-Flatin, J.P., Jelasity, M. and Zambonelli, F. (eds) (2007) 'Proceedings of the First International Conference on Self-Adaptive and Self-Organizining Systems (SASO2007)', July 2007, Boston: MIT, IEEE Press.

Doran, J. (1998) 'Simulating collective misbelief', *Journal of Artificial Societies and Social Simulation,* 1 (1).

Doran, J., Palmer, M., Gilbert, N., and Mellars, P. (1994) 'The EOS Project: modelling Upper Palaeolithic social change', in N. Gilbert and J. Doran (eds) *Simulating Societies: the computer simulation of social phenomena,* London: UCL Press.

Edmonds, B. (2004) *Using the Experimental Method to Produce Reliable Self-Organised Systems: engineering self-organising systems,* Berlin: Springer.

Edmonds, B. and Bryson, J. (2004) 'The insufficiency of formal design', in 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), 19–23 August 2004, New York: IEEE Computer Society Press.

Edmonds, B. and Hales, D. (2003) 'Replication, replication and replication—some hard lessons from model alignment', *Journal of Artificial Societies and Social Simulation,* 6 (4).

Epstein, J. and Axtell, R. (1996) *Growing Artificial Societies: social science from the bottom up,* London: MIT Press.

Galan, J. and Izquierdo, L. (2005) 'Appearances can be deceiving: lessons learned re-implementing Axelrod's "evolutionary approach to norms"', *Journal of Artificial Societies and Social Simulation,* 8 (3). Online. Available HTTP: <http://jasss.soc.surrey.ac.uk/8/3/2.html>.

Gilbert, N. and Conte, R. (eds) (1995) *Artificial Societies: the computer simulation of social life,* London: UCL Press.

Gilbert, N. and Troitzsch, K. (2005) *Simulation for the Social Scientist,* London: Open University Press.

Hales, D. (1998) 'Artificial societies, theory building and memetics', Proceedings of the 15th International Conference on Cybernetics, International Association for Cybernetics (IAC), Namur, Belgium.

Hales, D. (2000) 'Cooperation without space or memory: tags, groups and the prisoner's dilemma', in S. Moss and P. Davidsson (eds) *Multi-Agent-Based Simulation, LNAI 1979,* Berlin: Springer: 157–66.

Hales, D. (2005) 'Self-organising, open and cooperative P2P societies—from tags to networks', in *Proceedings of the 2nd Workshop on Engineering Self-Organising Applications,* LNCS 3464, Berlin: Springer, 123–37.

Hales, D. (2006) 'Emergent group-level selection in a peer-to-peer network', *Complexus,* 3: 108–18 (DOI:10.1159/000094193).

Hales, D. and Arteconi, S. (2006) 'SLACER: a self-organizing protocol for coordination in P2P networks', *IEEE Intelligent Systems,* 21 (2): 29–35.

Hales, D., Edmonds, B., Norling, E. and Rouchier, J. (eds) (2003) 'Multi-agent based simulation III', Proceedings of the 4th International Workshop, MABS 2003, Melbourne, Australia, July 2003, *Lecture Notes in Artificial Intelligence,* 2927, Berlin: Springer.

Hales, D., Marcozzi, A. and Cortese, G. (2007) 'Towards cooperative, self-organised replica management', Proceedings of the First International Conference on Self-Adaptive and Self-Organizing Systems (SASO2007), July 2007, Boston: MIT, IEEE Press.

Hales, D., Rouchier, J. and Edmonds, B. (eds) (2003) 'Special issue on model-2-model comparison', *Journal of Artificial Societies and Social Simulation,* 6 (4).

Halpin, B. (1999) 'Simulation in sociology', *American Behavioral Scientist,* 42 (10): 1488–1508.

Mollona, E. and Marcozzi, A. (2009) 'FirmNet: the scope of firms and the allocation of task in a knowledge–based economy', *Computational and Mathematical Organization Theory* (DOI: 10.1007/s10588–008–9049–8).

Moss, S. and Edmonds, B. (2005) 'Sociology and simulation: statistical and qualitative cross-validation', *American Journal of Sociology,* 110 (4): 1095–131.

Palla, G., Barabasi, A.L. and Vicsek, T. (2007) 'Quantifying social group evolution', *Nature,* 446: 664–7.

Polhill, G. and Izquierdo, L. (2005) 'Lessons learned from converting the artificial stock market to interval arithmetic', *Journal of Artificial Societies and Social Simulation,* 8 (2).

Polhill, G., Izquierdo, L. and Gotts, N. (2004) 'The ghost in the model (and other effects of floating point arithmetic', *Journal of Artificial Societies and Social Simulation,* 8 (1).

Popper, K. (1968) *The Logic of Scientific Discovery,* London: Hutchinson.

Ramanath, A. and Gilbert, N. (2004) 'The design of participatory agent-based social simulations', *Journal of Artificial Societies and Social Simulation,* 7 (4).

Rouchier, J., Cioffi-Revilla, C., Polhill, G. and Takadama, K. (2008) 'Progress in model—to-model analysis', *Journal of Artificial Societies and Social Simulation,* 11 (2).

Will, O. and Hegselmann, R. (2008) 'A replication that failed: on the computational model in "Michael W. Macy and Yoshimichi Sato: trust, cooperation and market formation in the U.S. and Japan", *Proceedings of the National Academy of Sciences,* May 2002', *Journal of Artificial Societies and Social Simulation,* 11 (3).