

Bootstrapping Knowledge *about* Social Phenomena *using* Simulation Models

Bruce Edmonds

Centre for Policy Modelling

Manchester Metropolitan University

Formidable difficulties face anyone trying to model social phenomena using a formal system, such as a computer program. The differences between formal systems and complex, multi-faceted and meaning-laden social systems are so fundamental that many will criticise *any* attempt to bridge this gap. Despite this, there are those who are so bullish about the project of social simulation that they appear to believe that simple computer models, that are also useful and reliable indicators of how aspects of society works, are not only possible but within our grasp. This paper seeks to pour water on such optimism but, on the other hand, show that useful computational models might be 'evolved'. In this way it is disagreeing with both naive positivist and relativistic post-modernist positions. However this will require a greater 'selective pressure' against models that are not grounded in evidence, 'floating models', and will result in a plethora of complex and context-specific models.

This paper takes a naturalistic and evolutionary view of science following such as (Toulmin 1967, Toulmin 1972, Popper 1972, Campbell 1974, and Hull 1988). However it differs from them as it does not claim an evolution of knowledge or ideas in the abstract, but rather an evolution of formal models (in this case social simulation models). I will not refer to this extensive literature since although close to what is being discussed below, is subtly oblique to it – I think the arguments will be clearer without that complication. This is a synthesis and development of many of my previous papers, especially (Edmonds 2000), so I apologise in advance for the number of self-citations.

Some of the difficulties facing social simulation

There are many difficulties facing the social scientist who wants to capture some aspects of observed social phenomena in a simulation model.

- Firstly, there is the sheer difference between the formal models (i.e. computer programs) that we are using as compared to the social world that we observe. The former are explicit, precise, with a formal grammar, predictable at the micro-level, reproducible and work in the same way (pretty well) regardless of the computational context. The later are vague, fluid, uncertain, flaky, implicit and imprecise – which often seems to work completely different in similar situations, and whose operation seems to rely on the rich interaction of meaning in a way that is sometimes explicable but almost never predictable. In particular the gap between essentially formal symbols with 'thin' meaning and the rich semantic associations of the observed social world (for example as expressed in natural language) is particularly stark. A gap so wide that some philosophers have declared it unbridgeable (e.g. Lincoln and Guba 1985).

- Secondly there is the sheer variability, complication and complexity of the social world. Social phenomena seem to be at least as complex as biological phenomena but without the central organising principle of evolution as specified in the neo-Darwinian Synthesis. If there are any general organising principles (and this is not obviously the case) then there are many of these, each with differing (and sometimes overlapping) domains of application.
- Then there is the sheer lack of adequate multifaceted data about social phenomena. Social simulators always seem to have to choose between longitudinal studies OR narrative data OR cross-sectional surveys OR time-series data never with the option of having all of these about a single social process or event. There is simply not the emphasis on data collection and measurement in the social sciences that there is in the 'hard' sciences, and certainly not the corresponding prestige for those who collect it or invent ways of doing so.
- There is the more mundane difficulty of building, checking, maintaining, and analysing simulations (Izquierdo et al, in press). Even the simplest simulations are beyond our complete understanding, indeed that is often why we need them, because there is no other practical way to find out the complex ramifications of a set of interacting agents. This presence of emergent outcomes in the simulations makes them very difficult to check. There is no feasible way to systematically check that our simulations in fact correspond to our intentions for them (in terms of design and implementation). The only ultimate solution is the independent replication of simulations – working from the specifications and checking their results at a high degree of accuracy (Axtel et al. 1996). However such replication is incredibly difficult and time-consuming even in relatively simple cases (Edmonds and Hales 2003).
- The penultimate difficulty that I will mention is that of the inevitability of background assumptions in all we do. There are always things to give meaning to and from the substrate of the foreground actions and causal chains that we observe. Many of these are not immediately apparent to us since they are part of the contexts we inhabit and so are not perceptually apparent. This is the same as other fields, indeed I argue elsewhere that the concept of causation only makes sense *within* a context (Edmonds 2008). However it does seem that context is more critical in the social world than others, since it can not only change the outcomes of events but their very meaning (and hence kind of social outcome). Whilst in other fields it might be acceptable to represent extra-contextual interferences as some kind of random distribution or process, this is often manifestly inadequate with social phenomena (Edmonds and Hales 2005).
- Finally there are the foreground assumptions in social simulation. Even when we are aware of all of the assumptions they are often either too numerous to include in a single model or else we simply lack *any* evidence as to what they should be. Thus there are many social simulation models which include some version of inference, learning, decision-making etc. within the agents of the model, even though we have no idea whether this corresponds to that used by the corresponding actors that they are supposed to represent. It seems that often it is simply hoped that these details will not matter in the end – a hope that is rarely checked and at least sometimes wrong (Edmonds 2001).

Despite this, there are those in the social simulation community that still hope that there will be simple simulation models that are reliable enough to be useful with respect to observations of social phenomena. Some of the reasons for this hope (or at least some of the stated justifications for it)

will be briefly discussed later. However, it should be clear that a naive positivist approach will face substantial difficulties at the very least. I will argue that there is a role for simple social simulation models as playthings with which to train our intuitions, but not a central role in the long-term development of models.

The role of formal simulation models in the social process of science

So given the difficulties above, how is it that simulation models could have a role in helping us understand social phenomena? I will illustrate this with an analogy, an analogy that goes back to (Popper 1972) and probably before – the analogy of biological evolution with that of the development of knowledge. It is important to realise that this is only an analogy and that knowledge *will* develop in different ways to that of biological species but an analogy which I will take a bit seriously in terms of the kind of consequences one might expect from such a process.

The point is that simulation models could form a role analogous to DNA in the evolutionary process. So that just as early organisms might have only been sufficiently adapted by chance to a relatively easy ecological niche, eventually they evolved to enable the effective exploitation of harsher niches. In the same way, although early simulation models might be almost completely inadequate for understanding particular social phenomena, across the community of modellers, trying variations on these models, they may evolve over time to be more effective for a variety of phenomena. The formality of the simulations is important because it is necessary to have precise and reliable replication of what is being evolved. The importance of a digital (as compared to analogue) genome is that it resists a drift in time – any mutation causes a discrete change which can be immediately selected out (unless exceptionally it was useful). Ideas, as such, are very important but are re-interpreted by each individual that hears them described by another – each time they are communicated they change a bit.

Some of the things that such an analogy suggest include:

- That the accurate replication of simulations is important;
- That the production of variations of simulations is important;
- That the kind of selection that the models undergo broadly determines the possibilities that the evolutionary process explores;
- That the resulting collection of models will be a ‘mess’, with huge families of different models and most models adapted to quite specific sets of social phenomena;
- That in the very long run there is not necessary any progress (i.e. development in any particular direction such as greater complexity), but rather adaption to the landscape;
- In the short and medium term there may be development in the sense of models being evolving that ‘fit’ a particular set of social phenomena better.

Thus (according to this picture) it is not, *so* much, the job of individual modellers to produce a simulation models that are completely adequate to particular sets of phenomena, but rather their task is to play their part in a wider process, *between modellers*. Below I discuss what characteristics each modelling step might play as a single step in this trans-model process.

One of the important characteristics of simulation models, especially individual- or agent-based simulation models is their expressiveness. That is they are able to more easily express a greater

variety of *structures* that might be useful against observed social phenomena than analytic formal models (at least as they are normally used¹). The greater ease with which structural variations can be tried and tested (compared to traditional analytic models) means that a greater variety of models with very different structures can be developed. To push the evolutionary analogy somewhat the present proliferation of simulation models could be thought of as being similar to the Cambrian explosion of multi-cellular organisms. There is a more important reason for the success of individual-based models in that they represent a step towards a more straight-forward, descriptive relationship between models and what is observed in that it is the essence of social phenomena that society *is* composed of a number of interacting entities (Edmonds 2003). This allows for more opportunity to bring evidence to bear upon them, a theme I will take up later.

Thus, in this picture, a substantial part of the importance of formal simulation models is as a consistent and replicable referents into which knowledge about social phenomena can be developed and encoded. Their existence allows for a stable and reliable reference to be established which aids constructive discussion of the phenomena. Indeed one of the advantages of building simulation models is that it often suggests questions for further empirical investigation (Moss and Edmonds 2005b). Contrast a science such as physics where it is the formal models that take centre stage and the ideas are treated more as guides to the models with, say, philosophy where a large part of the arguments are about which meaning or referent is being used (or should be used) in the discussion. It is not obvious that, for all its self questioning and sophistication that philosophy is more successful than physics.

If this picture is at all right, it has a major consequence for the practice of social simulation, in that the kind of selective pressure determines how such an evolutionary process develops. So that if models are selected primarily as a result of how much fun they are to play with, broadly that is what the whole social process will produce – if, on the other hand, models are selected with respect to their fit with observed evidence then there is a greater chance the results will be adapted to *that* selective pressure. This is likely to be the case even if the ‘environment’ that the evolutionary process is adapting to is complex, fractured and has lots of very specific ‘niches’. Thus this brings forward the question of what *sorts* of models and modelling might be helpful if the end goal of the whole social process is to understand observed social phenomena.

There are various analyses of exactly what is necessary for an evolutionary process to occur (see Hull 1988). That the conditions for the occurrence of a process that *is a member of an abstract class of evolutionary-type processes*, not for the particular process that is biological evolution. These sets of criteria differ from each other but they all include mechanisms or processes for:

- (Mostly) faithful spread of copies the core representation (**Reproduction**);
- Producing variations of these representations (**Variation**);
- Systematically selecting among representations, e.g. related to their success in some environment (**Selection**).

Sometimes these are further analysed, so selection might be separated into evaluation and differential survival, or combined so that success at reproduction *is* the differential selection. Nonetheless this particular division is sufficient for our needs here.

Clearly there is no problem with the processes of variation in social simulation at the moment. There are a so many models and simulations that the problem is more to judge how they relate than that they are too similarⁱⁱ. How variations are made is not so important, indeed some think that the variations *should* be independent of the selective pressure (e.g. Campbell 1960).

In a sense reproduction is trivial, given that we are using computer programs as our core representation. One just acquires a copy of the program and runs or inspects it. However, understanding one's own program is difficult (Izquierdo et al, in press) but understanding someone else's is even harder. Adequate documentation and access to source code and indicative results has been difficult – the documentation in the average academic paper of even the simplest model is clearly not enough (Edmonds and Hales 2003). However these difficulties can be ameliorated by having accessible model archives along with minimum standards of documentation (Polhill et al 2008). There are clearly issues concerning the reproduction of models, however these are largely technical in nature apart from establishing the norm that models need to be independently replicated if they are to be trusted (Edmonds and Hales 2003). Thus this paper will concentrate on what can be seen as the weakest link in this process: the *selection* of models within the community of social simulators since, given that there is effective reproduction and variation, the selection pressure on the models (which models are forgotten, which taken up for further investigation) will determine the direction and results of this process.

The Modelling Process

Clearly people construct computational simulations for all kinds of different reasons. Thus a simulation may have been designed for the purpose of: illustration, entertainment, aesthetic appeal, intervention, checking analytic results, exploring the properties of a simulation model, understanding some observed social phenomena (Epstein 2008 gives a list of 16 uses). Clearly models are attractive to other academics for a similarly wide range of reasons. Here I am considering what of these kinds are likely to be conducive to an evolutionary process which will ultimately further the understanding of observed social phenomena.

Here it is useful to think of modelling activities in terms of a picture of the stages involved. A classic view of this is shown in Figure 1 following, among others (Hesse 1963, Rosen 1985). Here an observed process is related to the inference in a formal model via the model set-up (classically this was the initialisation of the model derived from measurement of the target process) and the mapping back to the target process (which was classically a prediction). The arrows show the direction of the inference and not necessarily the direction of the use of a model so, for example, a model might be used to support or produce an *explanation* of the process by working backwards from the observed outcomes of the model backwards to a model that fits these results to an explanation in terms of the model set-up (the mapping into the model).

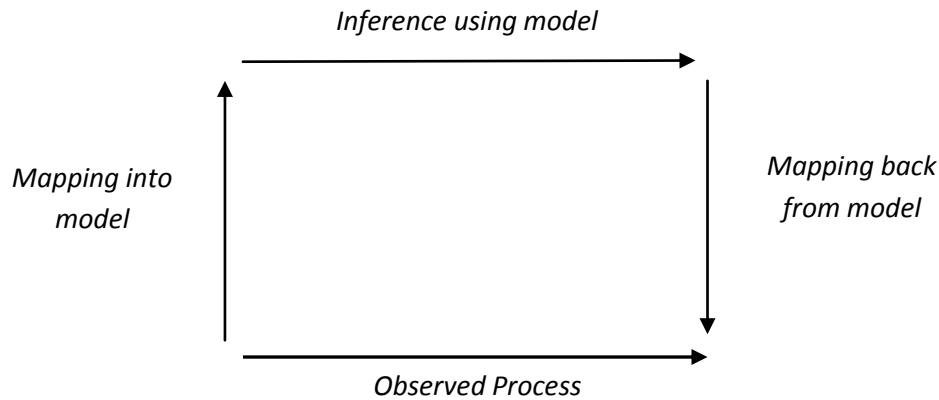


Figure 1. An illustration of the Modelling Relation following such as (Hesse 1963, Rosen 1985)

What is clear is that all three stages are necessary if one is to 'say' anything about the observed process. Often not all of these steps are described explicitly for example when the interpretation of the model is taken as self-evident. Thus, ultimately, the usefulness of a model in terms of understanding an observed process comes from how strong *all three modelling stages* are. Sometimes attention is focused on strengthening one of these stages but without regard to what effect this might have on the strength of whole chain. So for example, sometimes a simpler model might be used (e.g. using a set of solvable equations) so as to strengthen the inference stage, but at the cost of making the mapping between observation and the model tenuous because the model requires assumptions that are unlikely to hold in any case it will be applied to. Of course, it is very difficult to find a model that is simple enough to be useful but 'close' enough to the phenomena to say something about it – this is the hard job of science. However, it is simply mistaken to assume that a more tractable, understandable model (e.g. an analytically solvable one) is more scientific than a more complex one without taking into the resulting strength of all of the modelling stages taken together.

Another way in which a model is 'distanced' from the phenomena which it is supposed to be about is when the model is *not* mapped to and from anything observed but rather to a mental 'picture' of what is observed. Thus what often happens in social simulation is that a modeller has a conception of how a process works (or might work) and it is this conception that is modelled rather than anything that is observed (Edmonds 2001)ⁱⁱⁱ.

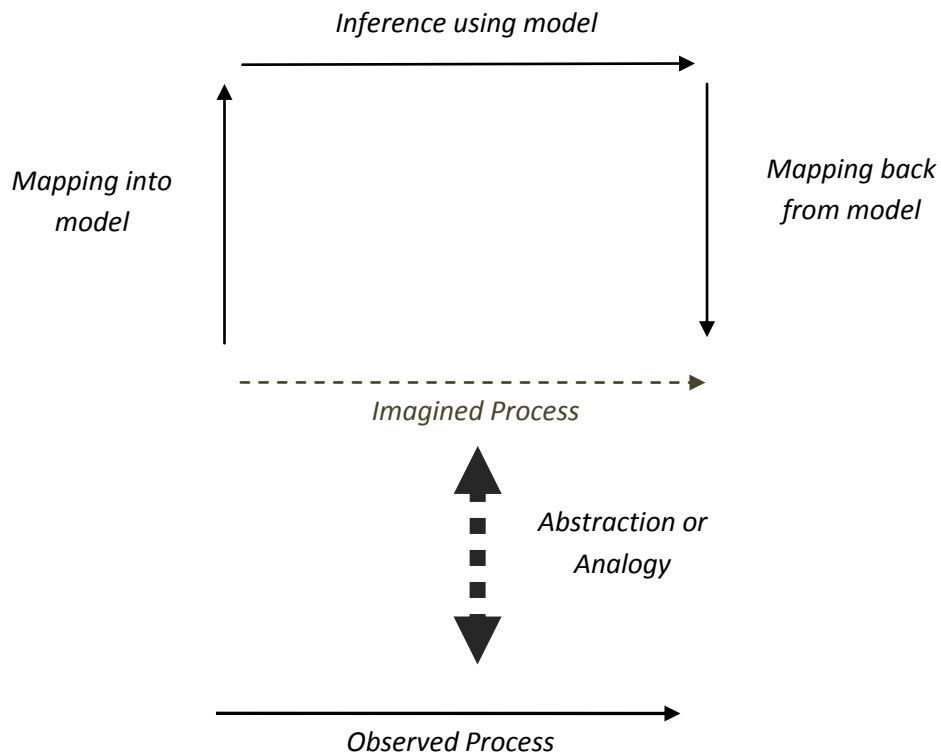


Figure 2. Modelling a mental conception (or analogy of an observed process)

Unfortunately it is often the case that the author's mental picture is conflated with the observed process and/or the connection between the concept and the observations left implicit. However in this case, how closely the modelling process can be informative is limited by how close the conceptual picture is to the phenomena itself. In a sense what one has in this case is an articulated analogy. This is an issue that will be taken up below.

Bootstrapping progress in model specificity

An objection to the picture of model evolution, painted above, might be that it does not fundamentally matter that models are disconnected to the evidence but more of articulated analogy because one can not escape simplifications and assumptions in any modelling. This view, which may derive from the post-modernism that pervades our culture, is that all such models are anyway merely a way of helping us think about what we see, and that we should be relaxed about how we used models to do this. The inability to completely escape assumptions in our models is used as a justification for an extreme relativism. To a person who thinks this way, it is entirely legitimate to use a model to encourage others to frame their thinking of the focus issues in a particular way, regardless of the extent it has been validated against evidence^{iv}.

However such arguments ignore the fact that the extent to which a model is constrained (or conversely determined by assumption) by the evidence is a matter of degree. If the result of confronting a model with evidence results in certain possibilities being excluded (or even just that certain possibilities are more likely) then this *is* some progress. That the evidence does not constrain the possibilities to a *unique* model/process does not invalidate it as science, it just means that the modelling process gave us less than the maximum information. Further, this is some progress even if this is conditional upon a large number of assumptions. As long as each modelling step is

sufficiently close to the evidence that some more possibilities are indicated or excluded and/or the assumptions behind the model revealed and documented then it is possible that gradually assumptions can be discovered and possibilities narrowed until the results are useful. In other words it is possible to gradually *bootstrap* increasingly accurate models in terms of the possible trajectories^v they suggest and an increasing understanding of the conditions under which one can rely on such models.

To see that such a bootstrapping process is possible, consider the closely analogous case of the development of measurement. For in the case of measurement, *even though* any measurement relies on other measurements and on reliable theory about the measurement process for its development, it *has* been possible to develop increasingly accurate methods of measurement over time. The process of achieving these levels of measurement has been successful, starting with inaccurate methods whose scope was uncertain (that is the conditions which they relied upon to give reliable answers were largely unknown). So techniques of measurement that only gave approximate results were used to help discover or test theories that allow better measurement, and techniques that only work in very special circumstances (i.e. rely on many assumptions about the conditions of the measurement process) were used to calibrate other techniques that might measure something else more accurately.

Similarly, if our modelling activities are appropriately directed then a slow bootstrapping of social simulation models is possible. It may be that each model *does* rely on a large number of assumptions but if one of the following holds it can still be a vehicle for modelling progress (as part of a community-wide evolutionary process).

These are that the assumption have been:

- partially validated in other models;
- the scope of how the model is applied has been established as one where the assumption is safe;
- or the nature of the assumption has been sufficiently investigated so as to understand that its impact is negligible or what the effect of it not holding might be.

This is not to say that such an evolutionary process will always be globally progressive, in the sense of slowly approximating an identifiable truth because social phenomena itself may well change over time and certainly the goals for modelling it will. However, it does mean that it is possible that locally this can be progressive in the sense the models can build upon previous models and refine their results.

This is in contrast to the re-use of models which are not constrained by evidence but rather are expressions of conceptions of processes, since with this later type it is relatively easy to produce a new model for each viewpoint or idea about a process or set of processes. Thus such modelling of conceptions will tend to rather increase the number of possibilities rather than hone them down because an analogy is not a strong selector on models (compared to evidence) due to its elastic nature. Also given that our models *do* have a strong influence on how we think about what we observe and think about then using ideas as the primary source of selection upon our models risks circularity and thus spuriously reinforcing what we think we know.

Modelling activities that promote model evolution

Some kinds of modelling that can play a part in such a bootstrapping of knowledge are described in this section. Of course this does not mean that other kinds are illegitimate (e.g. the others listed in (Epstein 2008)), just that the ones highlighted can play a part in the long-term evolution of model development.

Evidence-driven model selection

The first of these is where a model is developed that is strongly and identifiably constrained by some evidence. That is, the reason the model is published is due to it being successful (or more successful than other models) at explaining the evidence. Other ways of saying this are, that more of the evidence is consistent with the possible processes that the model reveals, or that the model (being consistent with the same evidence) rules out some processes or indicates some are more likely than others. This does not require the model to be consistent with *all* the available evidence, but just that it *contributes* to the understanding of the match between models and the evidence in identifiable cases. Nor does it require that models are very general (covering a range of cases), they can be conditional and specific.

The kinds of evidence involved can be very varied. Indeed it is one of the most significant advantages of agent-based simulation is that it allows a broader range of evidence to be applied: both narrative and quantitative (Moss and Edmonds 200?). Whilst it is true that some kinds of evidence are susceptible to various biases, it is fundamental to science that it does not ignore evidence without a very good reason – that it does not currently fit a particular modelling technique is not such a reason (but a reason to change the technique).

Thus I am disagreeing with (Moss 2008) where he states that the purpose of evidence-lead social simulation is “*intended precisely to represent the perceptions of stakeholders in order to bring clarity to scenarios built to explore the possibilities*”. Such an advantage is essentially ephemeral and could not play any part in a longer-scale inter-scientist process, since making the conceptions of the stakeholders (e.g. experts and/or participants) *precise*^{vi} could use *any* framework or model structure as long as it is formal. For such an exercise to be useful a model has to be more than precise, it has to be consistent with the knowledge of the stakeholders concerning the phenomena being studied (as well as any other evidence). It does not greatly matter that the judgement of the extent of the model coherency with the evidence is made by stakeholders rather than a professional modeller (each is susceptible to their own, albeit different biases^{vii}), but that the nature of the judgement and the extent of the agreement is laid bare and it results in a closer agreement of the model with the evidence (including crucially the opinions of the stakeholders). Thus the various participatory processes of model building (Barreteau et al 2003) *can* contribute to the development of social simulation modelling to the extent that evidence (including the stakeholder’s personal knowledge) is brought to bear on criticising the model design and outcomes. Clearly, according to the view of this paper, the *more* the evidence can be brought to bear and the more evidence that is brought to bear the better. Thus if it is possible to cross-validate a model in the sense of (Moss and Edmonds 2005a) that is better than only including using stakeholder opinion. However it is the total constraint of the evidence that counts, the opinions of many closely-involved stakeholders critiquing a model in terms of both design and outcomes may be a more effective constraint upon a model than an abstract time-series data set.

In order for an evidence-based modelling exercise to be most useful the assumptions and evidence on which it has been based or tested should be made as transparent as possible. In particular if any of the personal narrative evidence that is used is contested then this should be declared, so that the results of the modelling exercise can be seen as relative to the assumption that this evidence is reasonable.

A model can be ‘fitted’ with relative ease to either a set of known outcomes or a specification of the design. Constraining a model by only one of these is not enough for us to know that the match between a model and the evidence is worth communicating to others. In other words, only using one of the two mapping stages shown in Figure 1 is insufficient to significantly constrain what simulation trajectories can be produced. Two classical ways of using all the mapping stages is for the prediction of unknown data or the explanation of known data.

When predicting unknown data the simulation is initialised based on an observed case and the simulation is run *before* the simulator has seen the exact set of data it is predicting. This is a very strong test of a simulation outcome, judging only those that match a set of unknown data in the specified respects as acceptable. This is especially strong if the model predicts *unexpected* outcomes which then turn out to be correct. Prediction is such a strong test that even weak, propensity or negative prediction is some kind of real test of a simulation. Such weak tests include: that is a prediction that a particular thing (that might be expected) will not occur, that the outcomes will have a detectable propensity in a particular direction or that the outcomes will be of a certain well-defined kind. These are all predictions, albeit not precise predictions. If one has a model that predicts unknown outcomes then the outcome is strongly related to what is observed, even if the processes of the model that are not part of the prediction do not resemble that of the observed process.

Using a simulation to generate an explanation is far more common in social simulation. Here it is acceptable to fit *known* outcomes. The explanation represented by the simulation process gives an explanation of those outcomes in terms of the assumptions that are used to build the simulation. So if those assumptions are implausible then the explanation generated is equally implausible. Thus in this case it does matter if the processes of the model don’t strongly resemble that of the observed process because in case the process is (in a sense) the *result* of the exercise. Although not trivial, it is quite possible to fit known outcomes with a variety of model structures, so one needs the mapping between the model and target system structures to meaningfully constrain the simulation model.

One way of escaping the difficult task of finding a model that will strongly relate to the evidence is to mix up the criteria for a predictive and explanatory model. Thus many social simulation models (analytic and simulation) have a structure that is not strongly related to that of the observed process (e.g. using off-the-shelf learning algorithms or unrealistic assumptions of rationality) but then only demonstrate a fit to *known* data in terms of the outcomes (e.g. out-of-sample data). Clearly this is relatively easy to do since one has considerable freedom to fiddle with the structure and you can have as many goes at fitting the out-of-sample data as you want to before publishing the successful version. Such a model fails to be either a successful predictive or explanatory model and it does not represent a significant demonstration of a model which is constrained by the evidence. Sometimes this sort of model is combined with a use as an analogy to think with. This aspect is discussed further below.

Revealing model assumptions

In any model, even if much of the model structure is dictated by the evidence, there will still be many assumptions that are used in the simulation design, either because: they are unknown; they are simply assumed being a background assumption; they are necessary in order to get a simulation to run; or they were accidentally added into the model during its implementation. Some of these assumptions will not impact significantly upon the results – that is to say those aspects of the outcomes that are deemed to be significant in terms of what is observed will not be affected by them. However in complex simulations it is very difficult to tell which assumptions will be critical in this way and which will not. Further which assumptions are critical will, in general, depend on which aspects of the outcomes are deemed significant (in terms of the mapping back to the observed process). Their criticality will also depend on truth of other assumptions, so one assumption may not have a critical effect on the results, but only in the presence of another assumption (which may or may not be articulated).

Thus uncovering the assumptions a model depends upon (or is critically dependent upon for certain outcomes) is a complex and intricate matter. However due to the fact that any simulation will (implicitly or explicitly) rely on them makes it an important matter. So examining and analysing a model to find what its assumptions are is an exercise that is worthwhile. However, it is unlikely to deal with all assumptions so that detailing the situation it is being applied in is necessary as the aspects of the outcomes that are being focused upon. Assumptions are dependent upon the modelling context and thus an investigation of them can't be totally divorced from this context (Edmonds and Hales 2005).

Pseudo-mathematical modelling work

A last kind of modelling endeavour is a an attempt to simply understand the middle, inference step of the modelling process. This would be the simulation equivalent of pure mathematics, not concerned with the mapping of the model to anything observed. This is fine as long as it judged by the same criteria as those of good pure mathematics: soundness, generality/applicability and importance.

- Soundness means that others can rely on the results of the investigation without going through all the details as to what the original investigators did. This is extremely hard to do for all but the very simplest simulations and would, at the minimum, include independent replication.
- Generality means that the results of the investigation can be used by a reasonable number of others – in a simulation context it means that someone could tell something useful about the outcomes for their model using the results from the investigation.
- Importance means that the results show how information about a simulation can be transformed into a different but useful form – bridging the gap between different kinds of representation or reveal some hidden universality.

Taken together these criteria can be interpreted as seeking to ensure that, whilst the investigation is not directly useful, that it has a good chance of being useful in the future.

Clearly although such a kind of investigation is conceivable it involves painstaking work. Instead much work *suggests* something about other simulations but is not directly applicable. Again the

different kinds of modelling task are often mixed – so an investigation that does not really relate to any evidence but does not achieve the criteria of noteworthy mathematics above might be bolstered by an imprecise justification in terms of potential applications or vague analogy with some other models or conceptions. It seems to me that people look at the practice in sciences such as physics and draw the conclusion that the same kind of approach will work with the social sciences. In physics the micro-foundations (the behaviour of the bits) is often extremely well known so the challenge is how these might combine in complex systems. Also in physics it is relatively easy to get hold of data to test a model. The culture in physics makes it far more difficult to excuse a model that does not correspond to the evidence. Lastly, the relative simplicity of their phenomena allows for simpler, more tractable models, which allows their assumptions and structure to be better understood. Such a benign environment for modelling allows for a more relaxed style of modelling, that would not be effective in most social spheres where the above advantages do not hold.

‘Floating models’ and personal knowledge

Clearly there are models within social simulation that do not meet the criteria of any of the categories above. Such models might be justified vaguely with reference to some phenomena of interest, use many assumptions that are justified solely in terms of their surface plausibility to the modeller, that are fitted loosely to some known data for outcomes, but are not general enough to be considered as a sort of pseudo-mathematics. Such ‘floating models’ (Wartofsky 1979) are often closer to an expressions of the conceptions of the modeller rather than a model of anything observed – they are closest to an analogy – a computational analogy.

To be clear, I am not saying that such floating models are useless. Playing with such models can inform ones intuition about the possible processes involved. Just as with other analogies they can be a powerful tool for thought. In the case of such computational analogies they could allow people to hone their intuitions concerning some quite unpredictable and (otherwise) counter-intuitive outcomes resulting from complex processes. Thus (Moss, Artis and Omerod 1994) found that the main usefulness of national economic models was not for forecasting (they correctly forecast when nothing much changed but missed all the crucial turning points in the economy) but in the enriched understanding they gave those who used and maintained them.

However, such utility is largely a personal utility. That is, the understanding and usefulness is not something that is readily transferable to other people, and certainly not as part of locally progressive evolutionary process – they give personal rather than public knowledge. Rather such analogies come and go with the current culture – they are essentially transient entities. Just because one is excited about one’s own model and it has changed your own way of thinking does not mean that it will be useful to anyone else. Thus floating models *are* useful but mostly in this personal rather than in a scientific way.

Despite this, there are a lot of such models in the literature. Whilst at the early stages of a new field it might be more important to generate variety of models, so as to start off an evolutionary process, if there is not a selective pressure directed towards the evidence from observed social phenomena then it is likely that any such evolutionary process will not result in models adapted to that phenomena. There are various justifications for such models, which include those listed below. I have argued against these before, so will only summarise the arguments here.

- **Simplicity** – the claim that simpler models are more likely to be true (or truth-like etc.). This is complicated because there are lots of meanings attributed to the labels “simplicity” and “complexity”. Simpler models *are* easier to deal with and to a large extent we are forced to limit the complexity of our models due to cognitive, temporal and computational resource limitations. Also there is some justification to the formulation that elaborating unsuccessful models is not a good strategy (Edmonds and Moss 2005). However, there is simply no evidence that simpler models are better related to the evidence (Edmonds 2007b).
- **The Law of Large Numbers** – the assumption that the ‘noise’ will cancel out en masse (i.e. is random). This is the assumption that those details that are not captured by a model *will* cancel out as random noise does, give sufficient sample size or number of simulation runs etc. This makes the assumption that the parts of the outcomes that are not the identified “signal” are, in fact, random. It is true that we may use a (pseudo-)random process in our models to ‘stand in’ for such unmodelled aspects but that does not *make* them random (Edmonds 2005). In many social simulation models such “noise” is demonstrably not random (e.g. Edmonds 1999) so the assumption that it holds for social phenomena is questionable.
- **Abstraction** – that abstracting from detail will result in greater generality. Adding details into a model that come from a specific situation does make it less general. However the reverse is very far from the case – that abstracting a model will give it more generality – because there are many ways of abstracting and one does not know (1) whether a pattern *is* generalisable to other cases and (2) that one has made the right choices as to *what* to abstract. The problem is that with social phenomena that human behaviour is often highly context-dependent, and so many details of a situation may be necessary in order to set and determine that context (Kokinov and Grinberg 2001).
- **Plausibility** – that an academic’s intuitions are sufficient to ensure relevance. Using one’s intuitions about what to include in a model is inevitable and useful. However this is different from a justification of a model to a wider public in an academic paper – the justification that a model does not only personally inform the thinking of the modeller but is worthy of being part of the inter-modeller discourse. *Firstly*, such intuitions are only the weakest possible pressure from the observations to the model, being highly indirect and elastic. *Secondly*, the intuitions of an academic are highly influenced by the academic culture the modeller inhabits. It is that case that the intuitions of a particular field are self-reinforcing and that Kuhn’s “theoretical spectacles” (Kuhn 1962) shape the whole framework within which a modeller’s intuitions will be formed. Evidence can play the part of disrupting such shared assumptions where they are mistaken, to the extent they are allowed to. The less evidence has a role (or the more excuses for failure to relate to evidence are deemed acceptable) then the more likely it is for such shared assumptions to be entrenched. I would argue that some of the success of physics is due to its relative intolerance of models that do not correspond with the available evidence.
- **Data Fitting** – that the model outcomes vaguely match that of some data. As discussed above such a match is too weak to significantly constrain models, so could only be effective if there were many other evidence-led constraints upon it, e.g. that the model processes were also validated against the opinion of several stakeholders. A simple fit of known evidence even when combined with vague plausibility may be taken as an indication that a

model is worth investigating but that is a different matter from claiming that it is worth others' time in understanding it.

If the models that get replicated and cited in the literature are simple, fun models that are conceptually attractive but not evidence-related then that is what is likely to be the direction of the field as it unfolds. Thus if we are serious about understanding what the social phenomena we observe (in contrast to an emphasis as to how we think about them) then a scientific norm that floating models are not acceptable as the essence of a public communication (e.g. an academic paper) need to be established.

Some Caveats and Complexities

Clearly the inter-scientist process concerning the development of knowledge and its representations (in this case simulation models) is a highly complex and complicated process. The evolutionary picture of it is a simplification, there is no exact analogy with biological evolution but rather with a generalised process of evolution that one might see in models of that process (e.g. in the field of Evolutionary Computation).

One simplification in this picture is that a modeller may well draw on more than one source in terms of the next model they make. Thus although parts of models may be replicated they will not necessarily make a neat genetic tree, but a rather messy network. I don't think this effects the arguments presented too much. However if there is a culture so that each new model is essentially coded from scratch with only the *ideas* from others being used then it is not clear that there would be enough faithful replication of model parts so that an evolutionary process would result.

Secondly models clearly form an ecology of models. Models are not always judged on their own but as part of model clusters (Giere 1988). So that it may be that one model depends upon another, so that if that second model is discredited (hopefully with respect to data) then the former will be undermined too. Thus it may be that although models tend to be developed separately the selection process acts on little ecologies of models.

Conclusion – What needs to be done

This paper is optimistic in that it argues that the development of simulations that are well-adapted to the evidence is possible. It is pessimistic in that it suggests that this will be a lengthy process that is more likely to result in a plethora of complex, context-dependent and conditional models.

Clearly this paper calls for a norm to consider the public usefulness of models so that people who wish to present 'floating' models in academic papers feel they have to justify themselves. Not that there is a 'correct' kind of model, but that there are higher standards as to what modelling activities are worth while being promulgated. Whether this occurs depends on what the community decides is essential to social simulation, what its core is in the sense of (Lakatos 1970).

Secondly that to promote such a long-term process as described above that the standard of model documentation and archiving is greatly improved. At the moment it is a difficult and time-consuming process to replicate even the simplest of models (Edmonds and Hales 2003). Elsewhere Gary Polhill and I have written about some of these factors (Polhill and Edmonds 2007), but also

there is the Agent-based model archive at the Open Agent Based Modelling Consortium (<http://www.openabm.org/>), which not only allows one to archive models with them but also encourages better standards of model documentation – the ODD (Overview, Design Concepts, Details) protocol (Polhill, Parker and Grimm 2008).

Finally, if we believe that simulation can help us understand complex social phenomena, and that the science of social simulation is a complex social phenomena, then surely we must (eventually) seek to simulate this process, rather than only talk about it (as Gilbert 1997, Edmonds 2007a start to do).

Acknowledgements

Thanks are due to so many people that they are difficult to list. They must include: Scott Moss, and the other members of the Centre for Policy Modelling, participants of the second and third EPOS workshops, Luis Izquierdo, post-modernists in general, and Ronald Giere for his sensible and insightful philosophy.

References

Axelrod R., (1997) Advancing the Art of Simulation in the Social Sciences, R. Conte and R. Hegselmann and P. Terna (eds) *Simulating Social Phenomena*, Springer-Verlag. Selected Papers TBA, Berlin, pp 21-40.

Axtell, R., Axelrod R., J.M. Epstein and M.D. Cohen (1996), "Aligning Simulation Models: A Case Study and Results", *Computational and Mathematical Organization Theory* 1(2), pp. 123-141.

Campbell, D. T. (1960), "Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes." *Psychological Review*, 67/6: 380-400.

Campbell, D. T. (1974), "Evolutionary Epistemology." In *The philosophy of Karl R. Popper*, edited by P. A. Schilpp, LaSalle, IL: Open Court, pp. 412-463.

Carley, K. and Newell, A. (1994). The Nature of the Social Agent. *Journal of Mathematical Sociology*, 19:221-262.

Edmonds, B. (1999) Modelling Bounded Rationality In Agent-Based Simulations using the Evolution of Mental Models. In Brenner, T. (ed.), *Computational Techniques for Modelling Learning in Economics*, Kluwer, 305-332.

Edmonds, B. (2000) The Purpose and Place of Formal Systems in the Development of Science, CPM Report 00-75, MMU, UK. <http://cfpm.org/cpmrep75.html>.

Edmonds, B. (2001) The Use of Models - making MABS actually work. In Moss S and Davidsson P (Eds.) *Multi-Agent-Based Simulation*, Lecture Notes in Artificial Intelligence 1979: 15-32. Berlin: Springer-Verlag

Edmonds, B. (2003) Against: a priori theory For: descriptively adequate computational modelling, In *The Crisis in Economics: The Post-Autistic Economics Movement: The first 600 days*, Routledge, 175-179.

Edmonds, B. (2005) *The Nature of Noise*. CPM Report 05-156, MMU. (<http://cfpm.org/cpmrep156.html>).

Edmonds, B. (2007a) *Artificial Science - a Simulation to Study the Social Processes of Science*. In Edmonds, B., Hernandez, C. and Troitzsch, K. G. (eds.) *Social Simulation: Technologies, Advances and New Discoveries*. IGI Publishing, 61-67.

Edmonds, B. (2007b) *Simplicity is Not Truth-Indicative*. In Gershenson, C. et al. (2007) *Philosophy and Complexity*. World Scientific, 65-80.

Edmonds, B. (2007) *The Practical Modelling of Context-Dependent Causal Processes – A Recasting of Robert Rosen’s Thought*. *Chemistry and Biodiversity* (special issue on Robert Rosen), 4(1):2386-2395.

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). <http://jasss.soc.surrey.ac.uk/6/4/11.html>.

Edmonds, B. and Hales, D. (2005) *Computational Simulation as Theoretical Experiment*, *Journal of Mathematical Sociology* 29(3):209-232.

Edmonds, B. and Moss, S. (2001) *The Importance of Representing Cognitive Processes in Multi-Agent Models*, Invited paper at *Artificial Neural Networks - ICANN'2001*, Aug 21-25 2001, Vienna, Austria. Published in: Dorffner, G., Bischof, H. and Hornik, K. (eds.), *Lecture Notes in Computer Science*, 2130:759-766.

Edmonds, B. and Moss, S. (2005) *From KISS to KIDS – an ‘anti-simplistic’ modelling approach*. In P. Davidsson et al. (Eds.): *Multi Agent Based Simulation 2004*. Springer, *Lecture Notes in Artificial Intelligence*, 3415:130–144.

Epstein, J.M. (2008). 'Why Model?'. *Journal of Artificial Societies and Social Simulation* 11(4)12 <<http://jasss.soc.surrey.ac.uk/11/4/12.html>>.

Giere, R (1988) *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press, 1988, 313 pp.

Gilbert, N. (1997) *A simulation of the structure of academic science*. *SociologicalResearch Online*, 2(2), <http://www.socresonline.org.uk/2/2/3.html>

Hesse, M. B. (1963) *Models and Analogies in Science*. London: Sheed and Ward.

Hull, D. (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*, Chicago: The University of Chicago Press.

Kokinov, B., Grinberg, M.: Simulating context effects in problem solving with AMBR. In Akman, V., Bouquet, P., Thomason, R., Young, R., eds.: *Modelling and Using Context*. Number 2116 in LNAI, Springer (2001) 221–234

Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago: The University of Chicago Press.

Lakatos, I. and Musgrave, A. (eds.) (1970), *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press.

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Newbury Park, CA: Sage Publications.

Morgan, M. S. and Morrison, M. (1999) *Models as Mediators*. Cambridge University Press,

Moss, S. (2008) *Alternative Approaches to the Empirical Validation of Agent-Based Models*. *Journal of Artificial Societies and Social Simulation*, 11(1), 5. <http://jasss.soc.surrey.ac.uk/11/1/5.html>.

Moss, S. and Edmonds B (2005a) *Sociology and Simulation: - Statistical and Qualitative Cross-Validation*, *American Journal of Sociology*, 110(4), pp. 1095-1131.

Moss, S. and Edmonds, B. (2005b). *Towards Good Social Science*. *Journal of Artificial Societies and Social Simulation* 8(4).

Moss, Scott, Artis, M. and Ormerod, P. (1994), "A Smart Macroeconomic Forecasting System", *The Journal of Forecasting* 13, (3) 299-312.

Olivier Barreteau and others (2003) *Our Companion Modelling Approach*. *Journal of Artificial Societies and Social Simulation* vol. 6, no. 1 <<http://jasss.soc.surrey.ac.uk/6/2/1.html>>

Polhill, J. G. and Edmonds, B. (2007) *Open Access for Social Simulation*. *Journal of Artificial Societies and Social Simulation*, 10(3).

Polhill, J. G., Parker, D., Brown, D. and Grimm, V. (2008). 'Using the ODD Protocol for Describing Three Agent-Based Social Simulation Models of Land-Use Change'. *Journal of Artificial Societies and Social Simulation* 11(2)3 <<http://jasss.soc.surrey.ac.uk/11/2/3.html>>.

Popper, K. *Objective Knowledge: An Evolutionary Approach*, Routledge, 1972.

Rosen, R. (1985) *Anticipatory Systems*. New York: Pergamon.

Simon, H.A. (1986). *The failure of armchair economics [Interview]*. *Challenge*, 29(5), 18-25.

Toulmin, S. (1967), "The Evolutionary Development of Natural Science," *American Scientist*, 55: 4.

Toulmin, S. (1972), *Human Understanding: The Collective Use and Evolution of Concepts*, Princeton: Princeton University Press.

Wartofsky, M. W. ; 1979 *The Model Muddle: Proposals for an Immodest Realism (1966) in Models: Representation of the Scientific Understanding*. *Boston Studies in the Philosophy of Science* New York, N.Y. vol. 48, pp. 1-11

ⁱ There would be nothing to stop one having a set of complex and discontinuous equations for each individual which would be equivalent to any agent-based simulation, however one would then have to use and treat such a system in the same way as a simulation program, numerically calculating the results, debugging the equations etc. There would be no hope of analytic solutions or knowing the correctness of the equations.

ⁱⁱ The reader may contrast this with the situation in neo-classical economics where there is a tight orthodoxy on the types of models that are acceptable but where the selection by the evidence is, at best, weak.

ⁱⁱⁱ Unfortunately I was too polite in (Edmonds 2001), where I pointed out that much social simulation modelling was of a conception rather than about anything observed and, as a result, some authors have taken that article as a *justification* for that kind of modelling. I trust that this paper will redress that balance.

^{iv} An intermediate but somewhat vague position is that a model is a mediator between us and the evidence as in (Morgan and Morrison 1999).

^v A trajectory is a term for the course that a simulation traces out as it progresses.

^{vi} Here “precise” means that the model gives a very specific answer (e.g. 3.1415962kg), this does not imply *accuracy* which is how well the model outcomes match what it is modelling.

^{vii} Indeed an advantage of participatory approaches is that it democratises the modelling process.