# The Post-Truth Drift in Social Simulation

*Bruce Edmonds,*
*Centre for Policy Modelling*
*Manchester Metropolitan Univeristy*

*"It is better to be clearly wrong than vaguely right" – Pat Suppes*[1]

## Abstract

The paper identifies a danger in the field of social simulation a danger of using weasel words to give a false impression to the world about the achievements of our field. Whether this is intentional or unintentional, the effect might be to damage the reputation of the field and impair its development. At the root of this is a need for brutal honesty and openness, something that can be personally difficult and that needs social support. The paper considers some of the subtle ways that this kind of post-truth drift might occur, including: confusion/conflation of modelling purpose, wishing to justify pragmatic limitations in our work, falling back to unvalidated theory, confusing using a model for a way of looking at the world for something more reliable, and seeking protection from critique in vagueness. It calls on social simulation researchers to firmly reject such a drift.

## Introduction

"Post-truth" was nominated as their word of the year by the Oxford Dictionaries in 2016[2]. It is a kind of attitude to statements where the truth is considered less important than its emotional or political impact. It has been used to designate utterances that support the interests of a particular group (e.g. by attacking an opposing group), as in the UK Referendum Poll or the 2016 US Presidential Elections. This shift of emphasis from concerns about truth to those about power can be traced back to the critique of the, so-called, "Post modernists", such as Derrida who rightly pointed out that one has to understand the power relations behind statements to understand them and not just take them at face value. As critiques like this percolated into the public sphere it transformed into a critique of any expert opinion that was seen to contradict a desired political or economic direction of travel. This paper seeks to identify and resist a similar shift within social simulation.

Understanding social phenomena is hard. The temptation is thus to convince ourselves (and others) that we are making more progress than an objective analysis would support – in other words, to move the goal posts to make the game easier. In any difficult task it is sensible to have intermediate goals, which everybody involved understands. However, newcomers to a field might mistake these intermediate goals for the true goals and so, over time, a slippage occurs whereby only these easier, intermediate goals are considered in terms of any evaluation of work. This has happened in many areas of social science where, as the sheer complexity of our subject matter became apparent, more modest achievements are acceptable for papers to be published and projects to be funded. For example, in economics the "neo-classical" school followed Friedman (1953) when he argued that representing the structure of economic exchange was unnecessary in models as long as the model can predict, but then it quickly became the norm that it was enough to merely fit *known*, out-of-sample data, to get published.

This kind of slippage is sort-of OK if all the participants understand the meaning and significance of current goals within a field, but becomes more dangerous if these researchers interact with the public or policy-makers who do not understand such an "internal" framework. This can result in policy makers and funders over-estimating the claimed or planned achievements in a field and basing their decisions on this misunderstanding. It can also confuse newcomers to the field who may mistake these intermediate goals for the true ones[3].

The resulting disappointment by the wider public holds the danger of a disillusionment with the field as a whole, meaning that funding and interest could prematurely dry up before results, that are reliable enough for others to use, can be achieved. An example of this is in the use of the word "predict". Some modellers use this word to mean almost any calculation using a model, distinguishing it from

[1] This phrase is attributed to many people but Pat used the sentiment frequently in his lectures.
[2] https://en.oxforddictionaries.com/word-of-the-year/word-of-the-year-2016
[3] Inevitably, these goals that were considered as merely stepping stones then become *the* standard in a field.

“forecasting” which means predicting future or otherwise unknown data. However this is not always understood by others – if a modeller says “we can use this model to predict where riots might erupt”, then non-modellers will think it could be used by the police to know where riots are going to happen before they do with a high level of probability.

The early over-hyping of a field can cause an initial burst of enthusiasm (and funding) but then a period of disillusionment as it fails to deliver (

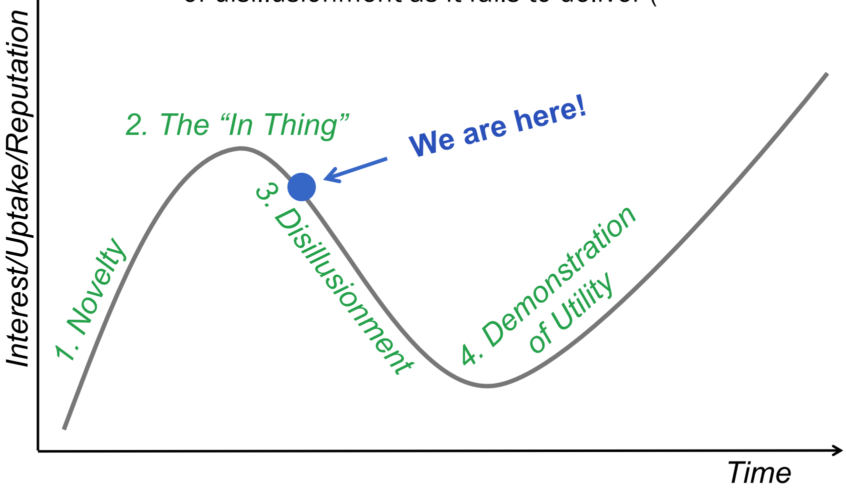


Figure 1). This was broadly the pattern for AI which claimed that useful intelligence would be achieved within a computer many decades before it started to appear. If one is lucky (as AI was once it had accepted the techniques of machine learning rather than restricting itself to reasoning and logic) one might survive to achieve demonstrable progress and start to gain acceptance again (but maybe under a different label). However, this upturn *can not be relied upon*! If academics refuse to adopt standards that matter outside their field and adapt their approach in order to meet these, there is the real possibility of being stuck in the “trough”. This seems to be the case with neo-classical economics which has resolutely refused to fundamentally question its assumptions or adopt rigorous and relevant criteria for judging its progress.

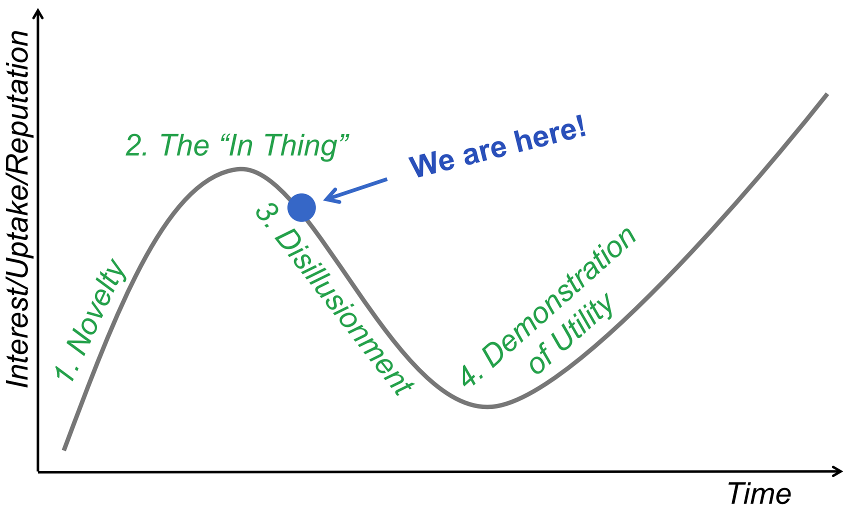


Figure 1. The “Hype Curve”

This is the choice facing social simulation – will it follow the comfortable path of weasel words and vacuous achievements, or will it “grasp the nettle” and attempt real progress in the knowledge that this may take much longer and that we may value intermediate stages less? Will it hold onto the ultimate importance of empirical truth or will it allow itself to focus on easier goals that give more of an appearance of progress? This paper discusses some of these trends in order to give them more prominence and so, maybe, cause researchers to think twice before following them.

## A Cautionary Example – The Club of Rome's "Limits to Growth" Model

In the early 1970's, on behalf of an international group under the name "The Club of Rome" a simulation study was published (Meadows et al. 1972) with the attempt to convince humankind that there were some serious issues facing it, in terms of a coming population, resource and pollution catastrophe. To do this they developed a system-dynamics model of the world. They chose a system-dynamics model because they felt they needed to capture some of the feedback cycles between the key factors – factors that would not come out in simple statistical projections of the available data. They developed this model and ran it, publishing the findings – a number of model-generated future scenarios – for a variety of settings and variations. The book ("Limits to Growth") considered the world as a single system, and postulated some relationships between a few macro variables, such as population, available resources, pollution etc. Based on these relations it simulated what might happen if the feedback between the various variables were allowed to occur. The results of the simulations were the curves that resulted from this model as the simulation was continued for future dates. The results implied that there was a coming critical point in time and that a lot of suffering would result, even if mankind managed to survive it.

The book had a considerable impact, firmly injecting the possibility that humankind could not simply continue to grow indefinitely into the public discourse. It also attracted considerable criticism (e.g. Cole et al. 1973) mainly based on the plausibility of the model's assumptions and the sensitivity of its results to those relationships. (For example it assumed that growth will be exponential and that delay loops are extended). Although the book did manage to put doubt about the wisdom of indefinite growth on the world agenda, it also had the consequence of bringing that kind of modelling into disrepute for decades.

The problem was that the book allowed an impression that the results of the simulations were *predictions* – a series of what-if scenarios which implied real forecasts concerning our joint future. Whilst they did add caveats and explore various possible versions of their model (depending on what connections there turned out to be in the world-system) the overall intent of the book was unmistakeable: that if we did not change what we were doing, by limiting our own economic consumption and population, disaster *would* result. This was a work firmly in the tradition of Malthus (1798) who, 175 years earlier, had predicted a constant state of near-starvation for much of the world based upon a consideration of the growth processes of population and agriculture.

The authors clearly hoped that by using a simulation (albeit a simplistic one by present standards) they would be able to make the potential feedback loops real to people. Thus, this was a use of simulation to *illustrate* an understanding that the authors of LTG had. However, the model was not presented as such, but as something more rigorous, more 'scientific'. A science-driven study that predicted such suffering was a more substantial challenge to those who thought growth could be allowed to continue indefinitely and, maybe for that reason, the book allowed that impression to be given.

The model was criticised on many different grounds, but the most effective was that the model was sensitive to the initial settings of some parameters (Vermeulen & de Jongh 1976). This raised the question of whether the model had to be finely tuned in order to get the behaviour claimed and thus, since the parameters were highly abstract and did not directly correspond to anything measurable, the applicability of the model to the world we live in was questioned. Its critics assumed that since this model did not, hence, produce reliable *predictions* that it could be safely ignored. It also engendered the general perception that predictive simulation models are not credible tools for understanding human socio-economic changes – especially for long-term analyses – and discouraged their use in supporting policy-making.

This paper is addressing the fear, that by over-hyping the success of our models and not being clear about the strength of their relationship to observed data, we may risk a similar disillusionment as occurred after the "Limits to Growth" was published.

## Confusion over Modelling Purpose

A model is not a picture of reality, but rather a tool for a particular purpose. The problem comes from there being many possible purposes for a model (Epstein 2008), but that many researchers in social simulation seem unable or unwilling to make their chosen purpose clear, with the result that their work is hard to judge.

Five such purposes are: prediction (in the forecasting sense), establishing an explanation, mapping out theoretical consequences, illustration, and as an analogy. Each of these is distinct and implies different approaches to development, checking, use, reliability etc. Establishing a model for one purpose does not justify it for another. If one avoids such clarity and allows for some vague mixture of purposes one risks confusion and unreliable science. If it is being suggested that a model can be used for a new purpose, it has to be justified for this new purpose. If a model is to have multiple purposes, it should be justified for each purpose separately. This is discussed in more detail in (Edmonds, forthcoming, 2017).

To drive home this point further, consider some common confusions of purpose to underline this danger.

1. Mapping Theory or Analogy → Explanation. Once one has immersed oneself in a model, there is a danger that the world looks like this model to its author. Here the temptation is to jump to an explanation of something in the world. A model can provide a way of looking at some phenomena, but just because one can view some phenomena in a particular way does not make it a good explanation.
2. Explanation → Prediction. A model that establishes an explanation traces a (complex) set of causal steps from the model set-up to outcomes that compare well with observed data. It is thus tempting to suggest that one can use this model to predict this observed data. However, establishing that a model is good for prediction requires its testing against unknown data many times – this goes way beyond what is needed to establish a candidate explanation for some phenomena.
3. Illustration → Understanding Theory. A neat illustration of an idea, suggests a mechanism. Thus the temptation is to use a model designed as an illustration or playful exploration as being sufficient for the purpose of a Understanding Theory. Understanding Theory involves the extensive testing of code to check the behaviour and the assumptions. An illustration, however suggestive, is not that rigorous. For example, it maybe that an illustrated process only appears under very particular circumstances, or it may be that the outcomes were due to aspects of the model that were thought to be unimportant. The work to rule out these kinds of possibility is what differentiates using a model as an illustration from modelling for Understanding Theory.

There is a natural progression in terms of purpose attempted as understanding develops: from illustration to description or understanding theory, from description to explanations and from explanations to prediction. However, each stage requires its own justification and probably a complete re-working of the model code for this new purpose. Whether people are not clear in their own minds about their purpose in modelling – whether they believe their model is multi-purpose, or whether its convenient to obscure – not being precise allows for a slippage between purposes and hence makes research harder to judge.

Some purposes for a model seem to be deliberately obscure or unclear. Looking to answer 'what if' questions does not specify the nature of the 'what if's that are being demonstrated. Does it mean that if something happens or is the case that one can predict the results? This is probably what might be assumed by policy makers if one told them your model answers such questions – *if the policy is thus then the outcomes will be thus*. It could be simply part of a theoretical exploration – *if we assume this then the outcome from my model would be this*. It could be just an illustration – *to show how if this happens then this could happen*. More likely it would be a kind of theoretical exploration of what might be the case if a certain explanation is true – *if this explanation is true, then if this is the case then this would occur as a result (if nothing else intervened)*. Thus saying one's model allows one to answer such questions leaves the exact role or reliability of the model, nicely obscure. A related example is that of scenario analysis/production which has a similarly large range of possible meanings.

In many publications and presentations, the authors are not clear about the purpose for their model. This may be because the model does not succeed under any heading, and so a lack of clarity about this is to their advantage. It may be because the modellers have just developed their model without thinking about its purpose. It may be that they are just confused about this issue. Whatever the case, if an audience can not clearly tell the heading under which they are supposed to judge a model, they should reject it.

## Not Facing Up to the Complexity of Our Subject Matter

The world (especially the social and ecological world) is not constructed for the convenience of us researchers. We cannot know, in advance of trying it, how many mechanisms and how much detail might be necessary to model any particular phenomena for any particular purpose. We are limited beings, with limited intelligence, imagination and time – there may be many things we can not directly understand and much detail we do not have time enough to include. This is not shameful – it is simply the truth and should be straightforwardly admitted.

However, there is a tendency to wish to justify our limitations on other grounds – that, somehow, these choices are the best ones, rather than pragmatic compromises. Some of these go as follows.

1. There is an optimum complexity for models. In (Grimm et al. 2005) it is suggested that there is an optimal complexity for models, saying "*We call this the ''Medawar zone'' because Medawar described a similar relation between the difficulty of a scientific problem and its payoff*" and citing (Loehle 1990). This was illustrated as in Figure 2. However, Loehle (1990) was arguing that this was the result of a pragmatic choice, following (Medawar 1967), concerning what kinds of problem might be worth tackling. A problem that is too easy might not impress ones peers enough to gain reputation or publication, a problem that is too difficult might not be solved. In other words to choose modelling targets that are doable within the resources available. What it does not say is anything about what level of complexity of model is best for any particular modelling target. It may well be that an good-enough model (for a particular case and purpose) is much more complicated that we have time or resources to explore. It maybe that an academic who has already determined what was going to be modelled to limit the complexity of their modelling, but this does not mean its optimal in any wider sense. A better choice would be to model something one has time to deal with adequately.

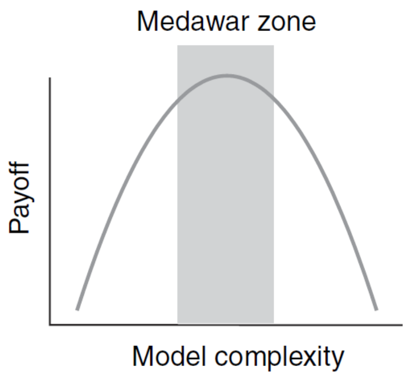


Figure 2. Implied Trade-off between model complexity and payoff from (Grimm et al. 2005)

2. Simpler models are more general. It is often implied, or worse merely assumed, that simpler models are more general (albeit at the cost of not being so accurate). The picture seems to be that simple models can get things roughly right, and increasing detail will make the model more accurate. To see the falsity of this, consider a linear equation model of some phenomena (e.g. the gas law equations), then "simplify" it by removing the variables and leaving a constant. The resulting model has less generality, not more[4], since it is now wrong almost everywhere. If you miss out a crucial component of a system in your model your model will not be fit for purpose – its generality radically *decreases*. The difficulty we face with social phenomena is that we have little idea which mechanisms or structures are essential – assuming we can just chuck out or ignore stuff without good grounds is wishful thinking – we may be lucky, but we can not rely on this.

[4] It is probably sufficiently accurate around one point only, if any.

3. **Complication can imply less complexity.** A dominant paradigm in the “sciences of complexity” has been to show that complex or new phenomena can be produced from the interaction of many relatively simple entities – that “*More is Different*” as Anderson (1972) put it. However, this is for illustrative purposes to make the emergence of complexity vivid – to prove that the complexity in the results emerged and did not result from the complication of the initial conditions. However, the difficulty in perceiving complexity in complicated systems does not mean that it is not there. There is no reason to suppose that adding complication into our models causes such complexity to disappear, just that it may then less clearly apparent (to our limited minds). It is theoretically possible that different processes or features may act to cancel each other out in very specific circumstances – but this would be sheer happenstance. If one varied the situation a little then there is no reason that they would continue to cancel. Thus diagrams such as Figure 3 taken from (Sun et al. 2016) and statements such as “the complexity of model behaviour may decrease after model complicatedness crosses a certain threshold” are, at best, highly specialised cases and, at worse, merely wishful thinking.

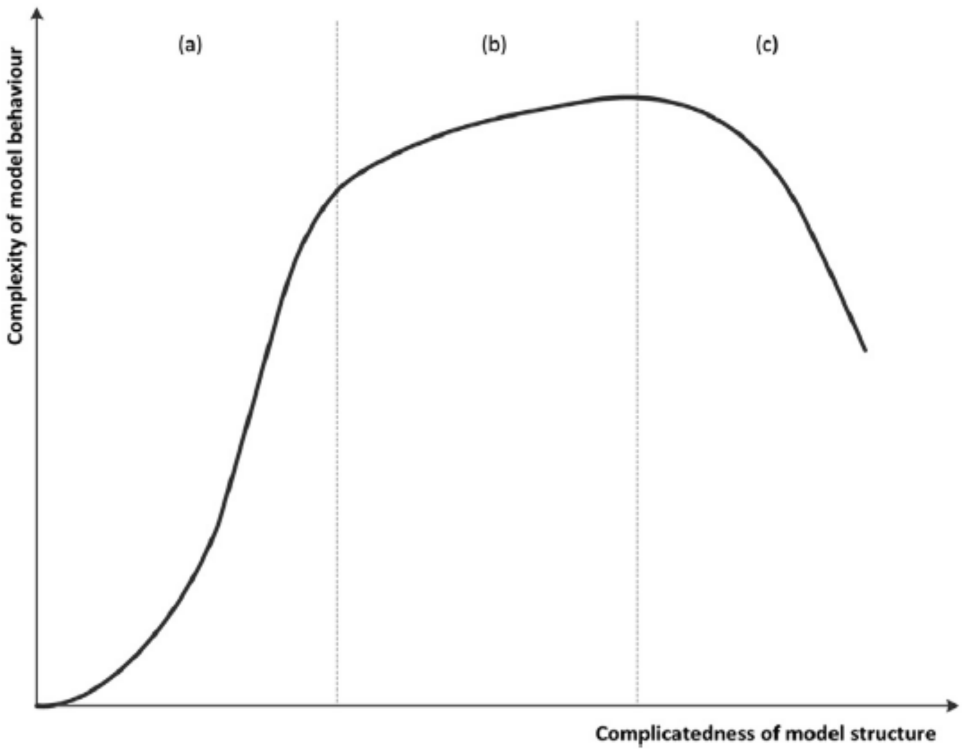


**Fig. 1.** Model complicatedness vs. Model complexity: (a) complexity increases exponentially with model complicatedness; (b) complexity increases at a lower ratio with model complicatedness; (c) complexity may decrease after certain threshold of model complicatedness.

Figure 3. Diagram from (Sun et al 2016) explaining how an increase in complication may result in a decrease in complexity after a certain level

Accepting these justifications at face value, may lead others to think that these reasons are real and hence limit their choices to the easy ones. They may be comforting but they also mislead! There are many *pragmatic* reasons for simplicity – simpler models are easier to build, to check, to communicate and to use. KISS is a sound engineering principle. However we cannot make assumptions about what level of complexity is best if we are trying to understand observed phenomena – missing out a key mechanism or detail can make a model fail[5]. The glib “For the sake of simplicity” [6] should be replaced by something more honest that admits to real, if human, limitations.

---

[5] Fail in its stated purpose

## Reliance on Existing Theory

One response to the difficulty of tasks such as predicting or explaining observed social phenomena is to retreat back to theory. In other words, to constrain what one researches only within the framework of an existing theory. This is a relatively safe route, since one can always say what one does is only critiquing or comparing existing theory. It also makes the job of a modeller *much* easier since it constrains the possibilities. Thus, assuming or staying within existing theory has clear pragmatic virtues for researchers. Some of the ways of doing this are as follows.

1. One might *compare* the possible consequences of Theory A vs. Theory B by implementing both of these in models for a specific context and see what results from each.
2. One might *assume* a certain theory, for example on human decision making, and create the rest of that simulation according to evidence etc.
3. One might explore the consequences of an existing theory in terms of some a specific implementation, mapping out the theory and its implications.
4. One might construct meta-theory to better understand the links and commonalities between existing theories.

These are all potentially valuable activities, but *only* if the theories are serious candidates for explaining or predicting observed phenomena. If the theories have no empirical validation then you might be simply wasting your time evaluating them and even worse if you are accepting them as assumptions as they may render your modelling irrelevant. The ultimate point of social simulation is observed social phenomena! The fact is that groups of academics have been convinced of un-validated theories and frameworks many times in the past only to later discover their irrelevance. The plausibility or history of a theory is not enough.

A common response to pointing out the simultaneous prevalence and empirical weakness of theory in the social sciences is that, without theory, generalisation is impossible. However, simply *wanting* generalisation is not enough, it is something that has to be achieved (usually with difficulty). It may well be that our theory is not strong enough to allow for generalisation, meaning we might have to be satisfied (at least for the moment) with more specific modelling exercises.

It is true that one cannot avoid theory altogether at any level. All complex simulation modelling will include many aspects that are not empirically suggested or justified, and these aspects may be reliable to different degrees. Some of these might be “suggested” by the literature, some by domain experts, and some might be common sense or guesses by the modeller. One has no choice but to be honest about ones assumptions and their reliability and declare these. However, the grander the theory and the less empirical support it has the more suspicion there should be. However one looks at this, the focus should be on the empirical chains of support, not the theoretical ones. A model or theory is only as strong as the empirical support it ultimately relies upon[7].

## Simulation Spectacles and Analogical Thinking

Kuhn documented how a theory can limit how one perceives the world, in particular to which kinds of empirical evidence one considers (Kuhn 1963), he called this wearing “theoretical spectacles”. There is a similar effect when developing a simulation model, but even stronger. Once one has worked on a simulation the whole world looks like one’s model for a long time afterwards. However this is also a danger, just because one can think of some phenomena using a model does not make this true[8].

The human mind is very good at using analogies to think about things (Hofstadter 1995). Almost anything can be used as an analogy, including simulations. Agent-based simulations seem to be particularly suggestive as to this, since one can envisage people in the place of the agents and can impute more into the simulation than is explicitly represented.

---

[6] For example, search the Journal of Artificial Societies and Social Simulation for the phrase “For the sake of simplicity”: http://jasss.soc.surrey.ac.uk/admin/searchresults.html?q=for%20the%20sake%20of%20simplicity
[7] Of course, that support could be indirect with empirical support for a model which supports a model which etc. but the end of the chain has to be empirical and the more un-validated assumptions that enter the chain the weaker it is.
[8] Or even useful for any particular purpose.

However, analogies do not tell you anything *reliable* about any situation or issue, they only provide a "way of thinking about things". There is a crucial difference between a model about the world and an analogy: an empirical model has a well-defined relationship with the world via data and specified measurements, with an analogy this relationship is re-invented every time it is applied. Thus analogies can be applied to many different situations but empirical models only to data and evidence.

Analogies are extremely useful things, they can provide us with suggestions and new insights, but this is not scientific knowledge – one can not rely on it because it has no empirical foundation. Analogies are useful as adjuncts to models – to think about the models and what we are doing with them – but it is the models themselves that form the backbone of science (Edmonds 2010). To make the difference clear I illustrate the comparisons involved below.

Common-sense tends to directly relate our understanding to what we observe without any intermediate stages (Figure 4). Scientific comparisons involve more intermediate stages, with data and formal models playing a significant part (Figure 5) – here the whole chain has to strong for the science to be reliable. Science has now a reasonable track record at producing reliable knowledge, even when this knowledge is surprising and initially clashes with our intuitive understanding. The kind of analogical reasoning using a model I am critiquing is where there is little or no comparison between the model and data, except maybe some indirect comparison via our intuitive understanding (Figure 6).

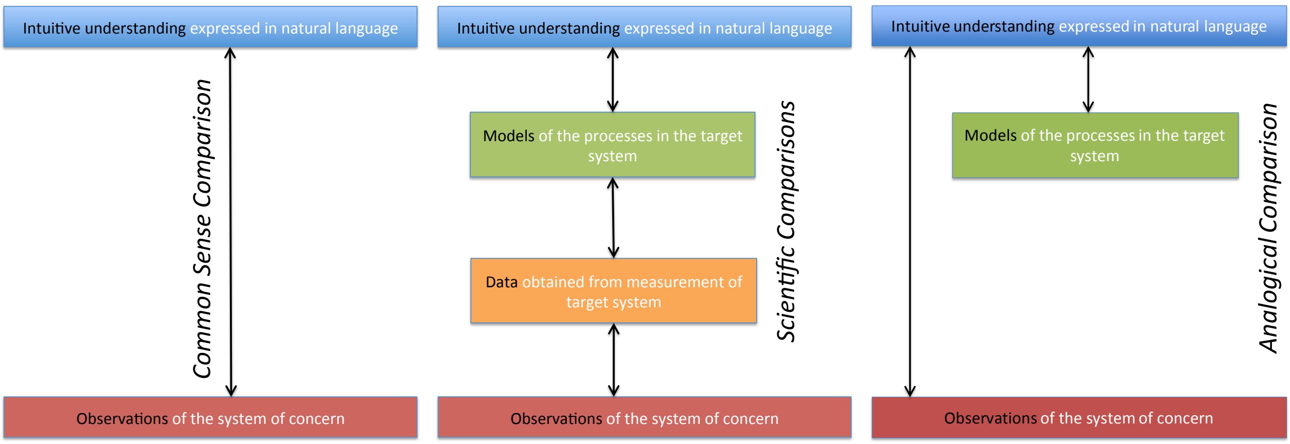


Figure 4. Common-Sense | Figure 5. Scientific | Figure 6. Analogical

There are some that would suggest that all models (including simulations) are "just" different ways of thinking about things, e.g. as in (Leydesdorff 2001). Saying "All models are wrong" is one way of expressing this[9]. However, taken seriously, this view would mean that there was no point in developing or checking models – the models would change as they were developed and checked but could not say they were "improved" at all by the effort of the modeller.

## Weasel words

To reiterate, there is *no* problem with analogical thinking – using a model as a way of thinking about things – as long as everybody is clear that this is *all* that is happening. Analogical thinking can give useful insights and ideas, it just does not give reliable ideas – they are not "scientific" in any sense that the public would understand[10]. Analogical thinking can happen in many ways and use different kinds of artefact, including the following.

- Developing and using conceptual or abstract formal models (mathematical or computational);
- Developing a model for explanatory purposes of some particular situation but only "validating" this against a set of intuitive understandings;
- Producing scenarios or comparing scenarios about complex future developments;
- Using a simulation as a means of communicating ideas, or mediating between stakeholders (as in participatory approaches to simulation).

[9] Which is, quite apart from anything else, irrelevant since models are tools not predicates.
[10] That is, analogies play a significant part in science, since ultimately we humans have to think about things, but its reliable conclusions have to be based on evidence so they can be relied upon.

What *is* a problem is when modeller gives a misleading impression – that such models and thinking have a more reliable basis without this being justified. This misleading impression may be due to living "within" a field too long, and mistaking its intermediate goals for ones the public would value. It may be due to pressure to claim significance for reasons of reputation, publication or career progression. It may be that they promised more in a grant proposal than they could deliver, and so seek to obscure the true significance of the results in project reports. However where such misleading occurs – whether on purpose, by omission or commission – this must be resisted. Effectively claiming (or merely allowing others to think) that more has been achieved than is justified is a selfish and short-term strategy. The danger is that the whole field will lose public trust and hence funding and respect in the longer term.

"Weasel words" come into being (or get used) when people want to give a false impression without actually lying – they obscure the truth but in a way that can be denied if pressed. They are particularly persistent when they enter the common usage of a group of people in order to justify their methods/results/etc. to themselves and others. The military have many such to talk about killing people, e.g. "collateral damage". Calling the results of internal model calculations "predictions" when they do not reliably anticipate aspects of the observed world seems to me to be one such usage – I have not come across anyone outside the world of modelling that would understand this in this purely internal sense.

## Participatory Approaches and Stakeholder Input

Not all modelling is designed to be objectively about something in the external world. Sometimes a simulation model is *meant* to be a medium through which stakeholders communicate and/or negotiate. In this case, the model is designed to reflect the beliefs and assumptions of these stakeholders or to be a way in which differing beliefs and assumptions can be compared or discussed. Such approaches include "group modelling" and companion modelling. Such approaches are completely legitimate uses of simulations, as long as they do not claim to be anything else.

Of course, the input of stakeholders can also be useful as a source of evidence about how something being modelled works, albeit a fallible one. Expert or stakeholder opinion is better than simply guessing at how something is or works, if one lacks more reliable kinds of evidence. Thus input from others could be an element in an empirical model with purpose other than that of communicating or negotiating. However, in this case, the model needs to be justified for the purpose it is declared as doing. If, for example, a model was specified purely in terms of the opinions of various people without independent checks then it should not be claimed as a reliable empirical model of a system. If the expert opinion is just a starting point, or concerning a minor element and there is a lot of other independent empirical checks on the outcomes then it might be acceptable. However, in any case, the source and reliability of any such assumptions or inputs needs to be honestly compared.

## Risking Being Individually Wrong

Popper (e.g. as in Popper 1963) argued for the overriding importance of possibly being wrong – what he called "refutability". Conceptual models, analogical thinking etc. are not refutable, even in principle. The sheer difficulty of our subject matter may incline researchers to be defensive, to prefer not to open their research up to being shown to be wrong and hence prefer models that are more analogical in their relationship with data. However, the whole direction of social simulation is to be less vague, to make precise (albeit complicated) models of social phenomena – models that can be inspected, compared and tested against data (at least in theory). If computational social simulation is to make a real contribution, then surely it is in the critiquability of its representations.

Being clear about the purpose for ones model, makes it easier to judge but it also encourages authors to be more realistic about their aims – not implying it achieves more than it does. If we are clear about what our research currently does, and how it should be judged this increases the chance of progress in the field. If we are going to succeed in our ambitious task of understanding societies better, then this will have to be a cooperative process, with understanding being slowly "boot strapped" thorough the shared development of our models between researchers and maybe over generations (Edmonds 2010). The key advantage of using precisely defined representations (our code and computations) is that they can be inspected, compared, tested and developed by other researchers. We will always need analogies to think about our code, but these will come and go, what will persist (if anything!) are the models.

Being open about one's model and its results is a corollary of this clarity. It should be standard that the model code and documentation should be openly available to others. The best and most permanent way to do this is via a public archive of models. The accompanying documentation is an important part of this – the model code by itself is very hard to understand and hence critique. The code should be given an appropriate license so that others have the rights to inspect, change and reuse the code – this allows for the model properties to be better explored and exploited. How to effectively share model code is discussed in (Polhill & Edmonds 2016).

This shift from over-valuing individual success with particular models to a more open and collective way of working is needed to increase the overall progress of the field. A greater degree of transparency and honesty will aid that process. This is why it is much better to be "*precisely wrong*" rather than "*vaguely right*". A model with known defects might be built upon or used as a null model for comparison. A "vague model" (i.e. one with an ill defined relationship with data – the interpretation being vague) is more difficult to critique because one has much more "wiggle room" of ways in which the model *might* be adequate.

## Staging Modelling

Although mine is pessimistic outlook as to the difficulties of the social simulation project, it is not without hope. Some of this hope comes from thinking about and taking care concerning the process of modelling itself. Using simulation models to understand complex social phenomena is an indirect manner of understanding; we cannot hold in our mind all the complex interactions that constitute social phenomena, but we can capture some of these within simulation models. We may not fully understand our own models, also having a level of detail beyond what our minds can encompass, but we can indefinitely inspect and experiment with them. In this way we can check them part by part, and formulate testable theories about our simulations and why we get the results we do. If (and only if) our simulations relate to the world in a reliable manner, we can use them to justify complex explanations and maybe even predictions about that world, even if the detail of these is beyond us. Thus we obtain useful, if indirect, understanding about social phenomena.

The point is that there is no reason why this process of indirect understanding can not be iterated, putting more stages between observations of the world and our understanding of that world. Data is already a model of our observations (Suppes 1969), our models combine parts of our understanding and allow us to compare them against data. It might be necessary to have more stages, such as models of our models to help us analyse and understand them (as in Lafuerza et al. 2016a), or even models of our models of our models, as in Lafuerza et al. 2016b). It might help to be more explicit in terms of representing "patterns" (Grimm et al. 2005) and how they relate to data. It would be useful to explicitly model which aspects of the output are deemed to be significant and how these might related to data, as in (Thorngate & Edmonds 2013).

The result would be a more complicated set of related models (more like that in Figure 7), each of which does a smaller abstraction step. It would require more work and a lot more checking. The understanding gained would be even more indirect, but it is possible.

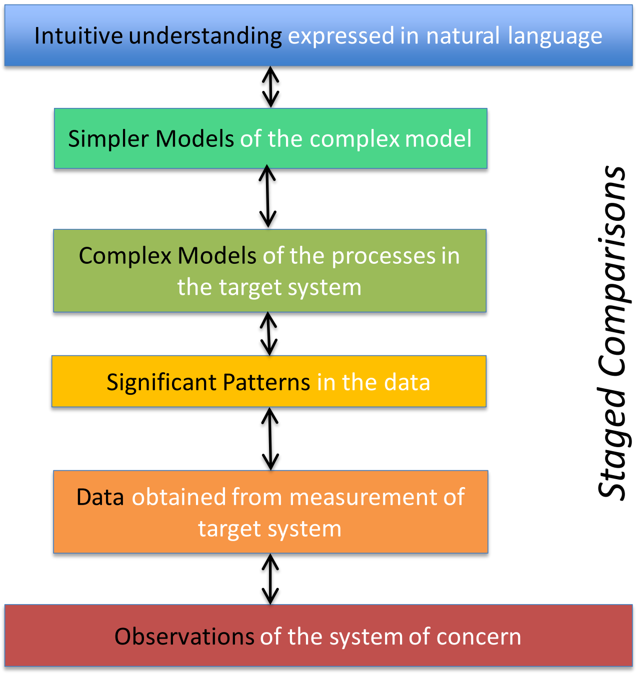


Figure 7. Staged Comparisons using many models

## Conclusions

There will be some that fundamentally disagree with the discussion above. There are some who see models as "just" another way of thinking about social phenomena, there are others that see models as "just" social constructions[11]. There is nothing wrong with this. What I am arguing against is any dishonesty (intended or unintended) – that there is any pretence that models that are *only* these things should be highly valued (by us, funding bodies, the public, policy-makers etc.). Our aims are very ambitious and making real progress towards them will be hard, but being more honest and modest is likely to result in more collective progress towards them. I call upon all social simulation researchers to reject and (politely) critique any such dishonesty – our collective reputation depends upon it.

## Acknowledgements

I would like to thank everybody who has argued with or against me on these issues. I have lost count of you all, but they include: Scott Moss, David Hales, Mike Bithell, Volker Grimm, Juliet Rouchier, Rosaria Conte, Alan McKane, and Guillaume Deffuant as well as all those who came to the workshop on validation held in Manchester in 2015. If you are not included in this list then you should have argued harder. I certify that my wife did *not* do any of the typing or checking of this paper[12]!

## References

ANDERSON, P. W. (1972). More is different. *Science*, 177(4047), 393-396.

COLE, H.S.D., Freeman, C., Jahoda, M. & Pavitt, K.L. (Eds.) (1973). *Models of Doom: A Critique of the Limits to Growth*. New York: Universe Books.

EDMONDS, B. (2010). Bootstrapping Knowledge About Social Phenomena Using Simulation Models. *Journal of Artificial Societies and Social Simulation*, 13, (1), 8, http://jasss.soc.surrey.ac.uk/13/1/8.html

EDMONDS, B. & Polhill, G. (2015) Open Modelling for Simulators. In Terán, O. and Aguilar, J. *(Eds.) Societal Benefits of Freely Accessible Technologies and Knowledge Resources* (pp. 237-254). IGI Global. (Open access version at: http://cfpm.org/discussionpapers/172).

[11] Though why this should stop them being objectively true or reliable, is never clear to me – I rely on socially constructed cars, trains and airplanes all the time!
[12] See the Twitter hashtag #Thanksfortyping

EDMONDS, B. (2017, forthcoming). Five Modelling Purposes. In Edmonds, B. & Meyer, R. (Eds.) *Simulating Social Complexity – a handbook*. Springer

EPSTEIN, JM. (2008). Why Model?. J*ournal of Artificial Societies and Social Simulation*, 11(4), 12, http://jasss.soc.surrey.ac.uk/11/4/12.html.

FRIEDMAN, M. (1953). *Essays in positive economics*. University of Chicago Press.

GRIMM, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., ... & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, 310, (5750), 987-991.

HOFSTADTER, D. R. (2008). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. Basic books.

KUHN, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago University Press.

LAFUERZA LF, Dyson L, Edmonds B, McKane AJ (2016b). Staged Models for Interdisciplinary Research. *PLoS ONE*, 11, (6), e0157261. DOI:10.1371/journal.pone.0157261

LAFUERZA, LF, Dyson, L, Edmonds, B & McKane, AJ (2016a). Simplification and analysis of a model of social interaction in voting, *European Physical Journal B*, 89, 159, DOI:10.1140/epjb/e2016-70062-2

LOEHLE, C. (1990). A guide to increased creativity in research: inspiration or perspiration? *Bioscience*, 40, (2), 123-129.

LEYDESDORFF, L. (2001). Technology and Culture: The Dissemination and the Potential 'Lock-in' of New Technologies. *Journal of Artificial Societies and Social Simulation*, 4(3), 5, http://jasss.soc.surrey.ac.uk/4/3/5.html

POPPER, K. R., & Hudson, G. E. (1963). *Conjectures and refutations*.

MALTHUS, T. (1798) *An Essay on the Principle of Population*. London: Johnson.

MEDAWAR, P. B. (1967). *The art of the soluble*. London: Methuen & Co. Ltd.

MEADOWS, D.H & Meadows, D., Randers, J. & Behrens, W.W.III (1972) *The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind*. New York: Universe Books.

THORNGATE, W. & Edmonds, B. (2013). Measuring simulation-observation fit: An introduction to ordinal pattern analysis. *Journal of Artificial Societies and Social Simulation*, 16(2), 14, http://jasss.soc.surrey.ac.uk/16/2/4.html

SUN, Z., Lorscheid, I., Millington, J. D., Lauf, S., Magliocca, N. R., Groeneveld, J., ... & Buchmann, C. M. (2016). Simple or complicated agent-based models? A complicated issue. Environmental Modelling & Software, 86, 56-67.

SUPPES, P. (1969). Models of data. In Studies in the Methodology and Foundations of Science (pp. 24-35). Springer Netherlands.

VERMEULEN, P.J. & de Jongh, D.C.J (1976). Parameter sensitivity of the 'Limits to Growth' world model. Applied Mathematical Modelling, 1, (1), 29-32.