

Using agent-based simulation to inform policy – what could *possibly* go wrong?

Bruce Edmonds and Lia ní Aodha

Centre for Policy Modelling
Manchester Metropolitan University
<http://cfpm.org>

“Naturally, politicians will look for any information or argument they can find to advance their agendas—that is their job” [1, p. 83]

Abstract. Scientific modelling can make things *worse*, as in the case of the North Atlantic Cod Fisheries Collapse. Some of these failures have been attributed to the simplicity of the models used compared to what they are trying to model. Agent-Based Modelling (ABM) pushes the boundaries of what can be simulated, prompting many to assume that it can usefully inform policy, even in the face of complexity. That said, ABM also brings with it new difficulties and potential confusions. This paper surveys some of the pitfalls that can arise when ABM analysts try to do this. Researchers who claim (or imply) that ABM can reliably predict complex phenomena are criticised in particular. However, an alternative is suggested – that of using ABM for a kind of uncertainty analysis – identifying some of the possible ways a policy can go wrong (or indeed go right). A fisheries example is given. This alternative may widen, rather than narrow, the range of evidence and possibilities that are considered, which could enrich the policy-making process. We call this *Reflexive Possibilistic Modelling*.

1. Introduction – the cautionary example of the North Atlantic Cod Fisheries Collapse

On the 2nd July 1992 Canada’s fisheries minister, John Crosbie, placed a moratorium on all cod fishing off the northeast coast of Newfoundland and Labrador. That day 30,000 people lost their jobs and hundreds of years fishing for cod off those coasts ended. The cod were declared commercially extinct [2]. Much research has been done seeking to explain the collapse, and many answers have been put forth – ranging from simplified narratives of overfishing, to environmental conditions, and poor management. A number of authors, however, have looked at the role played by fisheries science, arguing it and its models played a pivotal role in this collapse [2], [3].

Tracing the development of this fishery, and its management, the work of these scholars has recounted how, as the fishery was developed, both fisheries science and the policies it informed became increasingly centred on counting how many fish were in the sea [4] and predicting how many could be caught, with both of these processes feeding back on one another [2], [3], [5]. This, in turn, resulted in the development of increasingly intricate mathematical models which—partly due to the traditions in the field and the complex data they were trying to fit—became ever more divorced from reality [3], [5].

A series of scientific blunders (based on models) were made in the years leading up to the moratorium, in spite of repeated concerns being raised by inshore fishermen [3]. For example, a number of Commissions investigating the status of the cod, at the behest of the fishermen’s concerns, failed to make adequate inferences about the ailing stock health. The findings of these Commissions were buoyed by model-based assessments that depicted an increasing (rather than diminishing) resource base. These predictions saw development and investment in the fishery continue throughout much of the 1980s. During this time these figures were consistently challenged by inshore fishermen [3].

It was not until 1989 that what turned out to be an erroneous forecasting of fish stocks was corrected. The fisheries department issued its annual assessment based upon revised mathematical models, which indicated that abundance had been overestimated by as much as a factor of two [3]. The subsequent Harris Commission found that the fisheries department's estimates of stock strength were based upon data, methodologies, and models of such poor quality as to be essentially useless as a 'rational' basis for management or commercial planning. The executive summary of the Harris Report [6, p. 2] states that:

"...scientists, lulled by false data signals and, to some extent, overconfident of the validity of their predictions, failed to recognize the statistical inadequacies in their bulk biomass model and failed to properly acknowledge and recognize the high risk involved with state-of-stock advice based on relatively short and unreliable data series. Furthermore, the Panel is concerned that weaknesses in scientific management and the peer review process permitted this to happen."

Finlayson [3, pp. 12–15] argues that part of what had occurred was that scientists and policymakers had become so committed to their particular description of reality, and the idea of a strongly rebuilding stock, that this reality was read back into the scientific models.

Thus scientific models *can* make things worse – especially in cases where the modelling has been co-opted as part of the policy-making process! If this is true of the relatively simple but technical models that were being applied to questions of fisheries management in the 90's, how much more might this be true of the complex yet persuasive agent-based simulations we use now? This paper reviews some of the pitfalls of using Agent-Based Modelling (ABM) within the policy domain and ends by suggesting a positive way such simulations could be used, in a manner that may widen, rather than narrow, the range of evidence and possibilities that are taken into consideration.

2. Some pitfalls in using agent-based simulation for policy

There are a number of confusions and mistakes that can occur when analysts and policy actors interact. This is not very surprising because these roles come with very different sets of: skills, terminology, issues, assumptions and goals. Here we briefly discuss some of the pitfalls that may arise, from the point of view of a ABM analyst. We do not concentrate on those that are obvious to modellers of any kind, such as testing and checking your model thoroughly, but focus more on the pitfalls modellers may not have thought of, and those that are particularly relevant to ABM (though rarely are pitfalls exclusive). A fuller discussion of these may be found here [7] which this paper abstracts from and is concerned with a wider class of modelling techniques.

Institutionalised assumptions

All models, including complex models, contain assumptions about reality [1]. Oftentimes we make these choices without even realising it. These assumptions may be related to the theories we use, the fields we engage in, or wider cultural biases [8]. Whilst some may be based on empirical evidence, others might be wholly unverified [9]. They may, for example, be based on disciplinary theories that have never been proved beyond theory. Though we might not always be aware of these assumptions, together they will determine what goes into a model and what does not [10]. The danger here arises when erroneous assumptions that bear little or no resemblance to reality become institutionalised orthodoxies [11]. Models from neoclassical

economics and the untested models of human agency and behaviour embedded within these [12] are good examples of this. However, it has been argued that more complex simulation models are guilty of this also [13]. A good way to circumvent this problem is to be as explicit with ourselves (and others) about the assumptions we are making about reality, including those which may be taken-for-granted within (and often beyond) our fields, and be as comprehensive as possible. ABM does not necessarily require the kind of strong, theoretical assumptions that simple equation-based models did, but it often requires many *more* assumptions, albeit of a more mundane nature. If adopted without thinking any of these could be influencing the kind and range of results coming from the model (e.g. adding a 2D space to models).

Theoretical spectacles

Though we may not often acknowledge it, oftentimes our disciplinary biases [14], will likely entail a commitment to a worldview that leans towards some value system over others, or depicting aspects of that worldview over others in a model. Thomas Kuhn described this effect as wearing ‘theoretical spectacles’ [15]—the theories one believes leads one to only notice the aspects of the world that fit those theories. Empirical studies have suggested that when analysts spend much of their time producing and studying simulations rather than less mediated empirical representations (such as new kinds of evidence), critical distance becomes difficult to maintain [16]. Consequently, the danger of ‘theoretical spectacles’ is particularly acute for modellers, and can lead modellers to see the world ‘through’ their models, developing a strong confirmation bias [10]. The danger here is mistaking a model- or theory-driven view for reality, and hence failing to even consider alternatives. Although this is true of much modelling, it is particularly true of ABM, since ABM are often more descriptively accurate – and hence it is easier to project our views of reality on it. This becomes especially important in the policy-making process, as theory choices (and their underlying normative basis) can end up driving policy choices, with different theories likely implying very different policy prescriptions, often with quite diverging social visions and practical implications [17].

Model spread

One advantage of formal models is that they can be copied and used extensively with little effort. This is advantageous in the sense that it allows others to inspect, critique, and improve these models. It is disadvantageous, however, in the sense that once accepted, models tend to proliferate and spread uncritically. The ease of their re-use means that it can be tempting to reuse them, without retesting them, or evaluating their applicability to the task at hand. The danger here is that once a model becomes widespread others (both scientists and policymakers) then take this as a mark of its quality so it spreads even more. In an academic setting, this may not be too serious. However, in a policy setting the institutional acceptance of a poor model can have deleterious effects, with models working to concretise assumptions that though incorrect, may have moral weight [17]. The particular added danger with ABM models is that there are so many assumptions that it is difficult to tell when a model is being used beyond its original scope, since any one of the numerous assumptions it is based upon might be violated. Awareness of this pitfall is especially pertinent given studies have indicated that undue confidence can rise with distance from the site of knowledge production [16], [18].

Confusions over model purpose

Good models will have or should have, a clearly stated purpose. Such a model should have been designed with that purpose in mind and tested with respect to this. When published, its quality will have been judged against that purpose [19]. Models are not (or are almost never) general-purpose tools but more specific encapsulations of knowledge that have a quite specific scope of use. In many cases, if one does not know whether a model is being used beyond its scope, then it might be better to simply not use it at all—sometimes it is better to know the limitations of one’s knowledge than to think one has some idea (or baseline) of what is happening. This is a particular danger for ABM because it is a very flexible tool, that can be used for many different purposes [19], so it is easier to get these confused or make a wrong presumption about the intended purpose (especially if the authors are not clear about this).

Not knowing model limitations

Sometimes it is easy to forget the provisional relationship between our models and reality (18). It has been suggested elsewhere that this danger may be particularly acute with complex simulation models (such as many ABM), on account of the impression of verisimilitude that they give [16], [20]. All models however, have limitations. Limitations may arise, for instance, due to knowledge constraints with respect to the processes we are modelling, or computational constraints as to what can be feasibly modelled [1], [21], [22]. For example, some processes (e.g. political, cultural, and institutional) that are difficult to model oftentimes are not modelled because of these difficulties [20]. Alongside these, a model’s applicability is very likely to be contextually limited. In this regard, there are many examples of problems arising when ‘experts’ have applied explanations out of context [23], or when attempts have been made to make ‘predictions’ by drawing on observations under ‘similar’ conditions elsewhere [24]. Though some of these limitations are unavoidable, recognising them may help navigate other ABM pitfalls. Continually reviewing the suitability and usefulness of models (and the sub-models they contain), alongside running the model under different considerations e.g. its sensitivity analysis, may help circumvent the latter (more avoidable) limitations.

Uncertainty

There are multiple layers of uncertainty to be considered when modelling aspects of the world.¹ Some of these (though not all) will be related to our knowledge. For instance, there may be limitations with respect to our data, or inadequacies regarding our understanding of the processes being modelled [1], [16]. Alongside this, there will also be ‘model uncertainties’ relating to the question as to whether the model bears any relation to the world ‘out there’ [25]. In this sense, knowledge is always partial, and evidence is often open to various interpretations [26]. These uncertainties usually increase with the complexity of the system we are trying to understand, and the number of perspectives that are brought to bear on it. Failing to acknowledge these uncertainties, and partialities, or downplaying them—as is often the tendency in the public realm [20], [24]—is a further pitfall a modeller can fall into, and one that is likely to provide fuel for sceptics, reduce rather than expand the available options [27], and manifest later as ‘surprises’ [28].

Whilst (as we will argue later) that ABM can be useful for uncovering ways in which things may turn out different to expectations, this is still a danger if one is still viewing the model as

¹ Uncertainty is understood here as a situation where we don’t know what we don’t know [24],[52].

one where uncertainty is ‘contained’ or where one projects where models ‘behave’ upon what is being modelled.

Giving a false sense of security.

Sometimes introducing a model into the policy process, especially one that promises prediction, simply opens up space for a lack of action [1]. In this sense, models can lure us into a false sense of security and actually prevent us from doing anything useful, safe in the illusion that we can predict and hence manage future changes [24]. If we have a tool that can provide us with some forecasts, there is the risk of relying on this, and thus avoid responsibility for perhaps the worst-case scenario or the unknown scenario. In this respect, at both the modelling and policy levels, there is a tendency to focus on the best-case/most tolerable scenario and use a model to justify this restricted focus, rather than consider the full range of possibilities. Models in this regard can sometimes act as ‘symbolic substitutes’ for action, and work to gloss over ideological premises [29], providing instead reassurance for a particular worldview, alongside justification for a particular course of action [25].

Narrowing the evidential base.

As we saw with the case of the Newfoundland cod, although formal models can help us think about and understand things in a way we might not otherwise be able to do, they can also work to constrain the evidence base also, by side-lining other (perhaps less formal) forms of evidence. Sometimes other sources of evidence may be better suited to a task, than that which can be captured by a formal model—even one that is not constrained by the demands of analytical mathematics. In this sense, formalisation can lead to the exclusion of important nuances of a problem, oftentimes relating to processes that can be difficult to model, or value-based based issues on which there is no consensus. This might seem obvious, however, it becomes problematic in the sense that there is a tendency (discussed below) in policy-making to prioritise technical data (numerical data and the output of formal models) over other sources of evidence [30, p. 163]. It has been suggested elsewhere, that complex models may actually work to exacerbate this [31], which could work to further marginalise others forms of knowledge e.g. qualitative or local, place-based knowledge.

Inappropriate focus on ‘facts’ rather than values.

Over the past number of decades value based questions have increasingly been redefined as technical issues that can be solved through science [1], [26]. What this glosses over, however, is that knowledge (including scientific knowledge) is rarely completely independent of values [23]. Almost all models (even formal models) contain subjective judgements, and elements of normative social theorising [17] – especially when dealing with complex social or ecological situations. Whilst the use of expertise to inform policy in itself reflects certain values with respect to what constitutes evidence [23]. In this sense, the boundary between science and policy is fuzzier than is often acknowledged [1], [23]. Already existing research agendas complicate this further [12], as do the existence of multiple frames of analysis and thus interpretations of the problem, some of which will be incommensurable [29]. Though models can provide advice in this process, they cannot offer resolution to the value based disputes that define the political realm (which is a large part of the policy-making process) [26]. What often happens, however, is that a technical model—by focusing our attention on the

‘knowable’ [32]—can work to conceal these value-laden issues, and depoliticise this process [1] [33],[34].

Failing to understand the policy-making process.

There is a tendency (as discussed further below) to view the policy-making process as a linear process whereby a problem is (a) identified, (b) research is conducted, concrete evidence is supplied, which is (c) then used by policymakers to decide on the best way forward amidst competing interests. However, there is a lot of evidence to suggest that this is an oversimplified depiction of this process [1], [20], [26], [35]. Whilst experts do have a privileged role in forming policy, policies rarely flow directly from ‘facts’ [23]. In contrast, what often happens is that ‘facts’ can be used to legitimate particular courses of action [1] [26]. This is because the policy-making process is neither driven by ‘rational analysis’ or ‘expert judgement’ (nor should it be), but by ‘public’ debate over competing interests and values (i.e. this is democratic politics) [1], [26]. The goals of the policy might not be agreed on by all, and the extent to which the model complements or contradicts these goals will in large part determine whether it is accepted or rejected by policymakers [26]. For instance, some have suggested that the reason neoclassical economics has been so successful is that it provides an ideological function, rather than being empirically accurate [25].

A Summary of pitfalls

Pitfall	Mitigating Steps
Institutionalised assumptions.	Be explicit about assumptions, document them explicitly.
Theoretical spectacles.	Subject all assumptions to independent and reflexive examination.
Model spread.	Critically assess already existing models, even those that are taken-for-granted and used every day.
Confusions over model purpose.	Decide what the purpose of your model is – focus on this.
Not knowing model limitations.	Continually review the model, including the sub-models it contains. Do extensive model exploration and analysis.
Uncertainty.	Accept the uncertainty and be honest about it, in a manner that is not ambiguous.
Giving a false sense of security.	Don’t oversell your model and try to ensure caveats travel with the model and any conclusions drawn.
Narrowing the evidential base.	Bear in mind that not all knowledge can be formalised!
Inappropriate focus on ‘facts’ rather than values.	Is this simply an empirical question? Or is it also a normative one? (It is probably both).
Failing to understand the policy-making process.	How is my model <i>likely</i> to be used in this process – is the model suitable for that use?

3. The engineering/predictive paradigm for policy making

A common normative view of policy making (e.g. as described in the UK government's "Magenta Book" [36]) goes as follows:

1. Key performance indicators (KPIs) are designed to measure progress in obtaining policy goals (which are taken as given)
2. The possible policy options are described precisely
3. These are evaluated in terms of their *predicted* costs and benefits
4. The best policy is chosen and tried

This policy-choice process might then be embedded within an iterative approach which repeatedly (or occasionally) involves (a) evaluating current policies, (b) comparing the current policy against alternatives (using 2-4 above) and (c) changing or adapting the policy if this promises an improvement in outcomes.

Often, agent-based models are conceived of as helping policymakers in step 3 above, namely *predicting* the outcomes of different policies so that they can be systematically compared (usually using KPIs which can be estimated from model outcomes). Sometimes this is called "policy modelling", "model-based policy evaluation" or simply a "what if" analysis. I call this the "engineering paradigm" because it is analogous to the processes that are involved in engineering physical structures such as buildings or bridges.

ABM comes into the picture when it is perceived that simpler kinds of model are inadequate to doing the prediction that is needed in step 3 above. In other words, it is realised that the phenomena of concern is *complex*. Since ABM have the ability to represent sets of interacting agents, and explore the outcomes of different structures and behavioural rules, the hope is that such simulations will be able to provide the prediction that simpler models cannot. Thus, though the technology (and maybe the applicability) of ABM differs from previous modelling techniques, its role is the same.

However, it is unlikely that most ABM are able to fulfil this role. There are several reasons for this, including the following.

- Too many of the assumptions that go into the ABM are unreliable, for example aspects of agent decision-making in the simulation. Clearly, the more these aspects are based on good evidence, then this problem is reduced, however almost all ABM have so many assumptions that many unreliable (or even unquestioned) assumptions will remain.
- A fundamental kind of assumption is the list of those mechanisms that make it into the model and those that do not. There is no general meta-theory that tells us when we need any of a range of mechanisms (semantic communication, social norms, social influence, reputation, power, habit, practices etc. etc.). So it might be that if we miss one of these out that we totally alter the model results. A particular case of this is when we have not even thought of a mechanism or factor that turns out to be crucial later.
- The uncertainty in the target phenomena might be too great for prediction to be feasible. This may be due to a lack of adequate data or knowledge about key elements of the system, or it may be due to emergent chaotic outcomes that are produced by the system itself (however much one knows about it).

In practice, true prediction (that is, reliably anticipating data that is *unknown* to the modeller to a useful degree of accuracy) has turned out to be infeasible for most social phenomena². However, despite this deficiency (in terms of the predictive ability necessary to making this

² This does not mean it is impossible as [53] shows, or that we should not try to predict [54], but that it is only feasible in cases with limited set of outcomes and lots of data. There are no cases I know of where ABM have been predictively reliable.

paradigm *actually* work) does not stop researchers and analysts implying that they can do this. There seems a reluctance to tell policy actors (or grant-giving bodies) the truth – that such prediction will not be achieved³.

The point is that, if you cannot predict (and hence evaluate) the impact of policy, this paradigm does not work. This raises the question as to what the alternatives might be? In particular, whether there is any positive role for ABM in policy. In the next section, we propose one.

4. Using agent-based simulation for uncertainty analysis

Rather than predict what will happen, or even the probability distribution for what will happen, we suggest that a ABM is used to anticipate the various things that could happen. That is, move from a *probabilistic* approach to a *possibilistic* one. Assuming a ABM captures some of the complexity of what is being modelled, then some of the sequences of outcomes might be indicated in different runs of the simulation model. Note that such a simulation will not capture *all* of the possible outcomes – just some of them – thus it might not anticipate what *does* transpire. However, it could potentially map out more of the possibilities than those that could be envisaged without the simulation.

The usefulness of such an exercise is that it can inform a better *uncertainty analysis* of a situation or a policy. A risk-analysis estimates the probability of a known event occurring (usually something going wrong), an uncertainty analysis reveals some of the outcomes that one would not otherwise have planned for⁴. An analysis of the model runs should enable for the identification of a number of different kinds of outcome, each due to a different combination of processes.

Thus, whilst statistical and other methods might make predictions as to policy outcomes, a ABM might provide an accompanying analysis of the ways in which the policy might go unexpectedly wrong. Such an analysis can inform policy making and allow contingency planning if these outcomes are undesirable.

Once these possible outcomes have been identified, then the simulation can also be used to design ways of measuring the emergence of these outcomes. Equivalent measurement instruments can be implemented in the real situation to give the earliest possible warning of the emergence of deleterious outcomes. Such ‘early warning’ indicators and advance planning for these outcomes can allow policymakers to react in a more-timely manner, and hence mitigate the impact of these outcomes. In other words, such instruments can allow policymakers to better ‘steer’ or adapt policy.

Knowledge of a wider set of possibilities than those of just the status quo and the desired outcome can help inform the wider political debate. Thus instead of the success of a policy being the sole domain of a policymaker, a more informed discussion of the risks as well as the claimed benefits might be conducted by a more extensive set of stakeholders. Unlike with simpler models used for prediction, such an approach might widen, rather than narrow, the consideration of issues surrounding a proposed policy.

An example – fisheries

Fisheries management is an area where that pressure to predict has been particularly acute, and it has been argued that in the quest to do so ecological realism was sacrificed [2], [3],

³ Sometimes it is merely implied that prediction is possible using evasive or unclear language, e.g. meaning prediction of model results only, but allowing the stakeholder to think this means prediction of aspects of reality.

⁴ The risk vs. uncertainty distinction was originally made by Knight [55].

[37], [38]. Most fisheries management policies today are based upon, in one way or another, single stock assessments, calculations of maximum sustainable yield (MSY), and the allocation of quotas [4], [39]–[42]. What this necessitates practically is calculating the number of fish that are in the sea [4], and predicting the number of fish that will be produced each year by those that are left uncaught [43, p. 79]. Much of this assessment work is dominated by mathematical modelling of a particular kind.

This approach to nature has by now been widely critiqued with respect to its underpinning assumptions, and its practical failures. Built upon a number of partial theories (some of which are bio-economic hybrids combining both natural and economic ‘laws’ [44]), it assumes that nature can be understood by mechanistic laws, behaves deterministically, tends towards stability (and that humans have a role in maintaining this equilibrium), may be isolated and understood as single components, predicted and controlled, amongst others [2], [4], [5], [39], [43]–[46].⁵

Though we know today there are many problems with viewing nature in this manner [43], and sustained yield approaches to the management of natural resources (including fisheries) have garnered much critique [2], [4], [44]–[46] (precipitating shifts towards the development and application of alternative approaches to resource management), this approach to fisheries has become so institutionalized (both within fisheries science, and international policy), that by and large it continues to provide the basis for much policy in this area [2], [4], [5], [46]. As seen with the case of the Newfoundland cod stocks [3], [5], and others [4], [46], however, such an approach has oftentimes proven disastrous.

Thus, to illustrate the use of simulation for an uncertainty analysis, we applied an existing socio-ecological test bed [47] to look at the possible effects different fishing regimes may have on the underlying ecology. This is *not* to predict the effects of these regimes (even the average or expected outcomes), but look at the possible ways such regimes could go wrong.

Instead of modelling a few specific species, this test-bed allows for the evolution of a multi-species ecology which is constantly evolving and where the species are constantly interacting. We do not have space to explain this model in this short chapter, but refer readers to [47] for a more detailed explication.

In the version used here, plants and higher order trophic levels (fish) are distinguished – the plants compete for limited resources from the patches, and other individuals either eat the plants or each other for energy. The affordances between individuals and each other (or the patches they inhabit) follows [48] and has been shown (in other models) to result in the evolution of food webs with many of the characteristics of observed food webs [49]. The model consists of a 2D grid of well-mixed patches in which individuals can interact. This is illustrated in Figure 1.

⁵ Some of these assumptions, it has been argued, are more tightly related to the demands of mathematics, and the capacities of quantitative models, than anything else [38]. Alongside this, historians have been astute in highlighting that part of the reason this approach to fisheries gained such traction was that it provided an approach that was in line with a number of political and economic objectives [4].

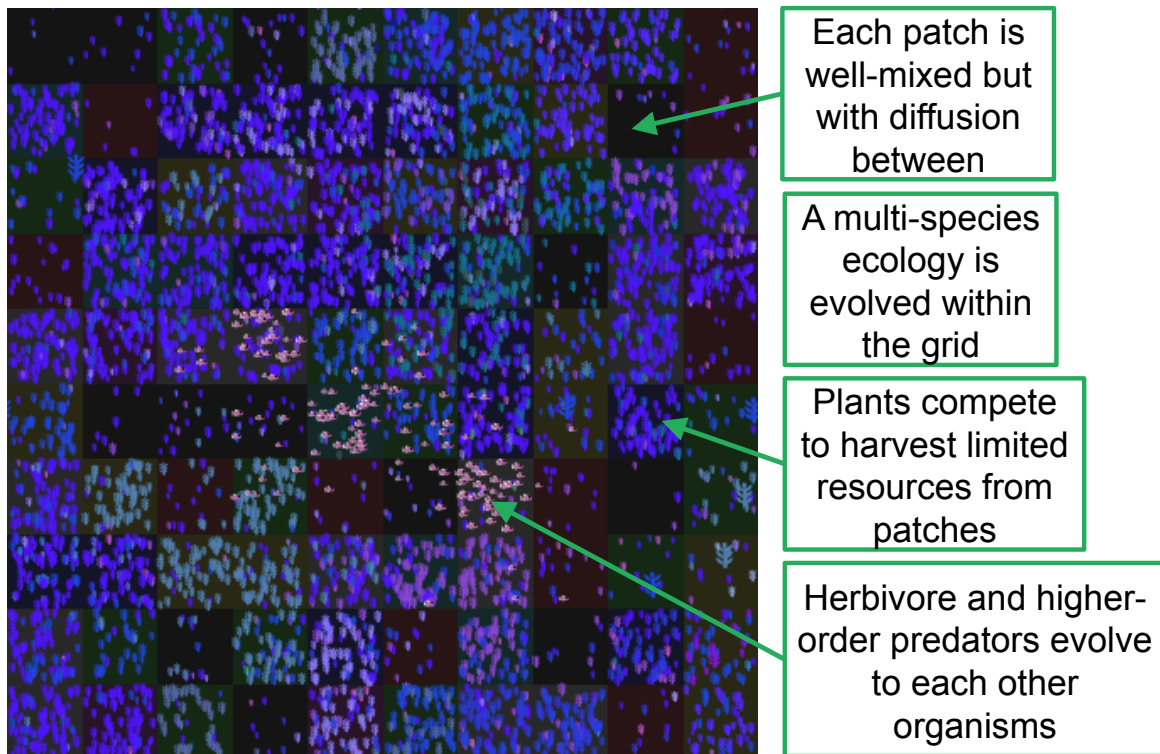


Figure 1. A multispecies, evolving fisheries model using a 2D grid of well-mixed patches

To produce a complex ecology of plant and fish species we evolve the whole ecosystem from scratch using an analogue of genetic evolution – starting with plants, then introducing herbivores and allowing higher predators to evolve from these. This complex ecology is then ‘saved’ for use as a consistent starting point for explorations of the impact of different catchment policies (or none). This is illustrated in Figure 2.




Figure 2. Evolving a complex ecosystem of species which is ‘saved’ as a consistent starting point for multiple runs under different extraction regimes (or none)

We then run the simulation starting in this exact state many times using different random seeds and applying different exogenous fishing extraction patterns to it. By examining the ‘trajectory’ of the runs over many ‘future’ simulation time intervals (ticks) we can see some of the possible outcomes of that policy and compare these against those of the null cases (no fishing).

To illustrate the model and its use, we run the model 20 times over 1000 more ticks at different levels of fish extraction, and look at the outcomes either over that time, or at the end of the runs.

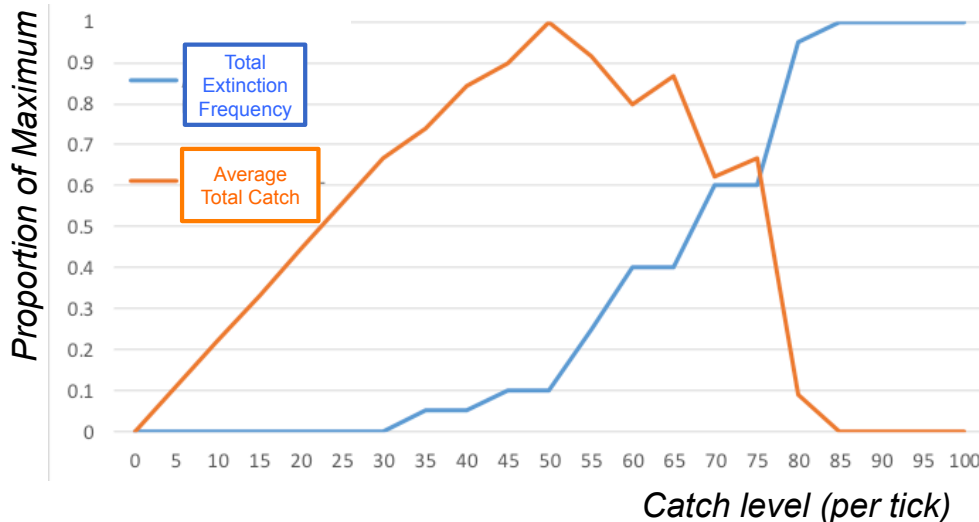


Figure 3. Average over last 100 ticks and 20 simulation runs for different levels of extraction of: (blue line) whether all species were extinct (orange line) total fish caught

Figure 3 shows a summary of the outcomes over the last 100 of those ticks (i.e. at the end of the run) summed over those runs and scaled by the maximum value that occurs. There the orange line shows a classic extraction curve with the greatest average catchment being at a level of 50 fish per tick caught. However the blue line shows the proportion of runs in which all fish had gone extinct by this time, and one can see that there was total extinction in at least one run at all levels of fish extraction above 30 fish / tick. In this model, at least, maximising “sustainable” fishing levels risks total extinction.

For the reasons why this apparent contradiction can occur, we need to look at some of the detail of the runs. One of the advantages of this approach is that we do not have to rely on average or equilibrium levels but can see some of the extremes that can occur dependent on the vagaries of chance and interspecies interaction.

Figure 4 shows the number of fish in each simulation over time, for two different levels of catchment (25 fish per tick and 35 fish per tick). Here one can see that although the fishing is done at a uniform rate, the number of fish varies chaotically all the time. This is true for this model even in cases where there is no fish extraction at all. What seems to be happening is that (a) fish numbers vary all the time up and down and (b) fish extraction creates a ‘danger zone’ at low fish levels – the higher the extraction rate the wider the zone. When fish levels wander into this zone rapid extinction of all fish occurs (as you can see at the right hand side of Figure 4 at 35 fish / tick extraction). Remember these are total fish statistics over all species present – if we looked into the variation of each species in each run, even more variation would be revealed.




Figure 4. Numbers of fish in each of 20 simulation runs with different levels of fish harvesting (each line is a different run)

Figure 5 shows the average number of fish left at the end of some longer runs (5000 ticks) for different catch levels. Here we see that, whilst above 60 fish / tick there is a catastrophic drop in fish numbers any level of fish extraction results in a lower level of fish (this is confirmed by more detailed series of runs done at low levels of fish extraction. That is, fish levels do not totally ‘recover’ when they are being constantly extracted. Similar results hold when the number of fish species are measured. In this model, at least, there is no threshold for fishing below which is safe – in any mature complex ecology there are always some species that are marginal and will be wiped out by even low levels of fishing. There is also, interestingly, no observable equilibria – if you continue running this model the lines do not ‘settle down’ but continue to gyrate.




Figure 5. Average number of fish at end of 5000 simulation ticks for different levels of catch with standard error bars (over 20 independent runs for each catch level)

Finally, we compare two methods of fish extraction: (a) a random uniform method of extraction (as used in all the above examples) and (b) a patch-by-patch method (where fishing is done by selecting a patch at random, taking all the fish there – or how much of the catch is left if less – then selecting another etc. until the quota is reached).

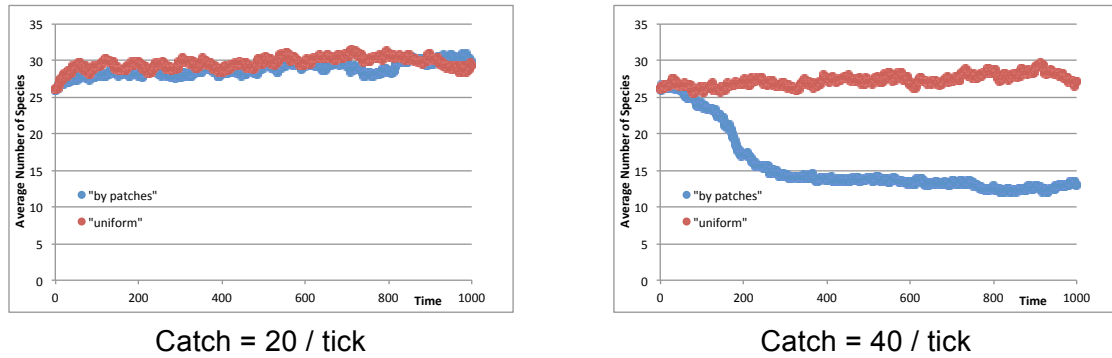


Figure 6. Average number of species over time and 20 runs, comparing a uniform, random method (red line) of catch compared to a patch-by-patch approach (blue line)

Figure 6 shows the average number of species over time for 20 runs over 1000 ticks for each of these options. In this version of the model, at lower levels of extraction each method has a similar impact on the number of species (on average), but at higher levels (40 fish / tick), there is a significant difference. Thus it could be that it is not just the *amount* of fish extracted that matters, but the *pattern* of that extraction.

We do not know if the complex ecologies being tested against are like the observed ones, these models do not *predict* the result of policy. But they do show plausible complex dynamics that cannot be ruled out in real, dynamic multispecies ecologies. They also illustrate what *might* happen when we move outside the realm of over-simple models and assumptions. Such an exercise could help warn us of some of the things that could go wrong in fisheries.

5. Concluding discussion

It is in the nature of many complex systems that we will never exhaustively know them [32], [50]. As we have seen, taking the road that pretends otherwise is fraught with dangers. Though it is unlikely that the partialities of our knowledge will be overcome, we have suggested a way that ABM might be used in a manner that could widen, rather than narrow, the consideration of issues surrounding a potential policy. To this end, we have proposed a process of *reflexive possibilistic modelling*. Rather than working to take the politics out of policy [1], such an approach, we hope, might provide a tool to assess future options under a range of scenarios, and from a range of perspectives.

That said, formal modelling approaches alone, even those that are reflexive and possibilistic, are unlikely to widen the possibilities within a policy-making process of limited reflexivity, where a multiplicity of perspectives are either not recognised or are systematically suppressed. Conversely, a process that encouraged open and transparent public debate with respect to the benefits, costs and uncertainties of different policy options (which would entail exposing the contingencies of our knowledge and the assumptions embedded within our methodologies) might, regardless of the quality of the modelling [17], [24], [35].

In this sense, taking multiplicity and plurality seriously means accepting that whilst models are useful, there is no one source of knowledge that can dictate the course of action under conditions of socio-environmental complexity [1]. In these instances different interpretations and interests will exist, and often the value frames of different actors in the policy-making process will be ‘incommensurable’ [51]. From this angle, forcing a consensual vision or model [1] would negate the importance of contestation for debating these values within the framework of democratic politics [51].

Acknowledgements

The authors acknowledge funding from the EU's Marie-Curie Horizon 2020 program as part of the Social Science Aspects of Fisheries for the 21st Century (SAF21) project, number 642080. We thank all those with whom we have had useful discussions on these subjects, including those at the University of Tromsø and at the ABM international workshop in Stockholm, July 2018.

References

- [1] D. Sarewitz, "Science and environmental policy: an excess of objectivity," *Earth matters earth Sci. Philos. claims community*, pp. 79–98, 2000.
- [2] D. Bavington, *Managed annihilation : an unnatural history of the Newfoundland cod collapse*. UBC Press, 2010.
- [3] A. C. Finlayson, *Fishing for truth : a sociological analysis of northern cod stock assessments from 1977 to 1990*. Institute of Social and Economic Research, Memorial University of Newfoundland, 1994.
- [4] C. Finley, *All the fish in the sea : maximum sustainable yield and the failure of fisheries management*. The University of Chicago Press, 2011.
- [5] D. Bavington, "Marine and Freshwater Fisheries in Canada: Uncertainties, Conflicts, and Hope on the Water," in *Resource and Environmental Management in Canada*, Oxford University Press Don Mills, 2015, pp. 221–245.
- [6] L. Harris, "Independent Review of the State of the Northern Cod Stock Executive summary and recommendations," 1990.
- [7] L. Aodha and B. Edmonds, "Some pitfalls to beware when applying models to issues of policy relevance," in *Simulating Social Complexity - a handbook*, no. 9783319669472, B. Edmonds, Ruth Meyer, Ed. Springer, 2017, pp. 801–822.
- [8] R. Conte, B. Edmonds, S. Moss, and R. K. Sawyer, "Sociology and Social Theory in Agent Based Social Simulation: A Symposium," *Comput. Math. Organ. Theory*, vol. 7, no. 3, pp. 183–205, 2001.
- [9] A. Satelli and S. Funtowicz, "When All Models Are Wrong," *Issues in Science and Technology*, vol. 30. University of Texas at Dallas, pp. 79–85, 2014.
- [10] J. D. Sterman, "All models are wrong: reflections on becoming a systems scientist," *Syst. Dyn. Rev.*, vol. 18, no. 4, pp. 501–531, 2002.
- [11] T. Forsyth, *Critical political ecology : the politics of environmental science*. Routledge, 2003.
- [12] S. Jasanoff, "(No?) Accounting for expertise," *Sci. Public Policy*, vol. 30, no. 3, pp. 157–162, 2003.
- [13] B. Epstein, "Agent-based modeling and the fallacies of individualism," in *Models Simulations and Representations*, Routledge, 2013, pp. 115–144.
- [14] P. D. Glynn, "Integrated Environmental Modelling: human decisions, human challenges," *Geol. Soc. London, Spec. Publ.*, vol. 408, no. 1, pp. 161–182, Jan. 2017.
- [15] T. S. Kuhn, *The Structure of Scientific Revolutions*. 1962.
- [16] M. Lahsen, "Seductive simulations? Uncertainty distribution around climate models," *Social Studies of Science*, vol. 35, no. 6. pp. 895–922, 2005.
- [17] R. Evans, "Economic Models and Policy Advice: Theory Choice or Moral Choice?," *Sci. Context*, vol. 12, no. 02, pp. 351–376, Jun. 1999.
- [18] B. Wynne, "Risk and Environment as Legitimatory Discourses of Technology: Reflexivity Inside Out?," *Curr. Sociol.*, vol. 50, no. 3, pp. 459–477, 2002.
- [19] B. Edmonds, "Different modelling purposes," in *Simulating Social Complexity - a handbook*, no. 9783319669472, Bruce Edmonds Ruth Meyer, Ed. Springer, 2017, pp. 39–58.
- [20] S. Shackley and B. Wynne, "Representing Uncertainty in Global Climate Change Science and Policy : Devices and Authority," *Sci. Technol. Human Values*, vol. 21, no. 3, pp. 275–302,

- 1996.
- [21] S. Sismondo, "Models, Simulations, and Their Objects," *Sci. Context*, vol. 12, no. 2, pp. 247–260, 2017.
- [22] E. Winsberg, "Sanctioning Models: The Epistemology of Simulation," *Sci. Context*, vol. 12, no. 02, 1999.
- [23] T. Forsyth, "Expertise needs transparency not blind trust: A deliberative approach to integrating science and social participation," *Crit. Policy Stud.*, vol. 5, no. 3, pp. 317–322, Oct. 2011.
- [24] B. Wynne, "Uncertainty and environmental learning. Reconceiving science and policy in the preventive paradigm," *Glob. Environ. Chang.*, vol. 2, no. 2, pp. 111–127, Jun. 1992.
- [25] J. R. Ravetz, "Economics as an elite folk science : the suppression of uncertainty," *J. Post Keynes. Econ.*, vol. 17, no. 2, pp. 165–184, 1994.
- [26] M. Heazle, *Uncertainty in policy making: Values and evidence in complex decisions*. Routledge, 2012.
- [27] A. Saltelli, P. B. Stark, W. Becker, and P. Stano, "Climate models as economic guides scientific challenge or quixotic quest?," *Issues Sci. Technol.*, vol. 31, no. 3, pp. 79–84, 2015.
- [28] L. Pellizzoni, "Knowledge, Uncertainty and the Transformation of the Public Sphere," *Eur. J. Soc. Theory*, vol. 6, no. 3, pp. 327–355, 2003.
- [29] J. S. Dryzek, "Policy Analysis and Planning: From Science to Argument," in *The Argumentative Turn in Policy Analysis and Planning*, Routledge, 1993, pp. 213–231.
- [30] W. Pearce, A. Wesselink, and H. Colebatch, "Evidence and meaning in policy making," *Evid. Policy A J. Res. Debate Pract.*, vol. 10, no. 2, pp. 161–165, May 2014.
- [31] A. Saltelli and M. Giampietro, "What is wrong with evidence based policy, and how can it be improved?," *Futures*, vol. 91, pp. 62–71, Aug. 2017.
- [32] S. Jasanoff, "Technologies of humility: researchers and policy-makers need ways for accommodating the partiality of scientific knowledge and for acting under the inevitable uncertainty it holds," *Nature*, vol. 450, no. 7166, pp. 33–34, Nov. 2007.
- [33] A. Benessia *et al.*, *The Rightful Place of Science: Science on the Verge*. 2016.
- [34] T. Forsyth, "Politicizing Environmental Science Does Not Mean Denying Climate Science Nor Endorsing It Without Question," *Glob. Environ. Polit.*, vol. 12, no. 2, pp. 18–23, May 2012.
- [35] S. Shackley and B. Wynne, "Integrating knowledges for climate change: Pyramids, nets and uncertainties," *Glob. Environ. Chang.*, vol. 5, no. 2, pp. 113–126, May 1995.
- [36] HM Treasury, *The Magenta Book: Guidance for evaluation*, no. April. 2011.
- [37] P. Holm, "Crossing the Border : On the Relationship BETween Science and Fishermen's Knowledge in a Resource Management Context," *Marit. Stud.*, vol. 2, no. 1, pp. 5–33, 2003.
- [38] R. C. Francis, "Fisheries science now and in the future: A personal view," *New Zeal. J. Mar. Freshw. Res.*, vol. 14, no. 1, pp. 95–100, Mar. 1980.
- [39] K. N. Nielsen and P. Holm, "A brief catalogue of failures: Framing evaluation and learning in fisheries resource management," *Mar. Policy*, vol. 31, no. 6, pp. 669–680, Nov. 2007.
- [40] T. Smith, *Scaling fisheries : the science of measuring the effects of fishing, 1855-1955*. Cambridge University Press, 1994.
- [41] L. Campling, E. Havice, and P. Mccall Howard, "The Political Economy and Ecology of Capture Fisheries: Market Dynamics, Resource Access and Relations of Exploitation and Resistance," *J. Agrar. Chang.*, vol. 12, no. 2–3, pp. 177–203, Apr. 2012.
- [42] G. M. Winder, "Introduction: Fisheries, Quota Management, Quota Transfer and Bio-economic Rationalization," in *Fisheries, Quota Management and Quota Transfer. Rationalization through Bio-economics.*, vol. 15, G. M. Winder, Ed. 2018, pp. 3–28.
- [43] S. Bocking, *Nature's experts : science, politics, and the environment*. Rutgers University Press, 2004.
- [44] J. Hubbard, "Fisheries Biology and the Dismal Science: Economists and the Rational Exploitation of Fisheries for Social Progress," in *Fisheries, Quota Management and Quota Transfer Rationalization through Bio-economics*, G. M. Winder, Ed. 2018, pp. 31–61.
- [45] J. Hubbard, "Mediating the North Atlantic Environment: Fisheries Biologists, Technology, and Marine Spaces," *Environ. Hist. Durh. N. C.*, vol. 18, no. 1, pp. 88–100, Jan. 2013.
- [46] O. H. Pilkey and L. Pilkey-Jarvis, *Useless Arithmetic: Why Environmental Scientists Can't*

- Predict the Future*. Columbia University Press, 2007.
- [47] B. Edmonds, “A Socio-Ecological Test Bed,” *Ecol. Complex.*, no. in press, 2018.
 - [48] G. Caldarelli, P. G. Higgs, and A. J. McKane, “Modelling coevolution in multispecies communities,” *J. Theor. Biol.*, 1998.
 - [49] A. J. McKane, “Evolving complex food webs,” in *European Physical Journal B*, 2004.
 - [50] A. Pickering, *The cybernetic brain : sketches of another future*. University of Chicago Press, 2010.
 - [51] B. S. Matulis and J. R. Moyer, “Beyond Inclusive Conservation: The Value of Pluralism, the Need for Agonism, and the Case for Social Instrumentalism,” *Conserv. Lett.*, vol. 10, no. 3, pp. 279–287, May 2017.
 - [52] L. Mehta, M. Leach, P. Newell, I. Scoones, K. Sivaramakrishnan, and S.-A. Way, “Exploring Understandings of Institutions and Uncertainty: New directions in natural Resource Management,” *IDS Discuss. Pap.* 372, vol. 1, pp. 1–48, 1999.
 - [53] N. Silver, *The signal and the noise : why so many predictions fail--but some don't*. Penguin Press, 2012.
 - [54] G. Polhill, “Why the social simulation community should tackle prediction,” *Rev. Artif. Soc. Soc. Simul.*, 2018.
 - [55] F. Knight, *Risk, uncertainty and profit*. New York: Hart, Schaffner and Marx, 1921.