Some pitfalls to beware when applying models to issues of policy relevance

Lia ní Aodha and Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan Univeristy

Abstract

This paper looks at some of the ways things can go wrong when mathematical or computational models are applied to inform policy on important issues. It looks at some of the pitfalls in the model construction and development phase, including: choosing assumptions, the effect of "theoretical spectacles", over-simplified models, not understanding model limitations, and not testing a model enough. It then goes on to discuss the pitfalls that can occur when a model is applied to inform policy, including: entrenched policies based on models with little or no evidential support, and how models can narrow the evidential base considered. It also looks at: confusions concerning model purpose and kinds of question they may answer, when models are used out of context, asking unreasonable things of models, when the uncertainties are too great, when models give a false sense of security, and when the focus should be on values rather than facts. This discussion is then illustrated with two examples, one economic and one from fisheries. It concludes that most of these problems stem from the interface between the modelling and policy worlds. It ends with some simple recommendations to reduce these mistakes.

1 Introduction

We all use models all the time, albeit usually informal mental models, but sometimes mathematical or computational models. These models help us think about situations we encounter – both familiar and unfamiliar. While such models, from the very informal to the most formal, can be helpful, they can also work to limit our understanding – biasing and even constraining how we think about things, or how we *might* think about things. An important characteristic of models is that they are simplified descriptions – representations that are designed by humans. As such, it is worth remembering that, in much the same way that certain arguments or concepts can work to obscure more than they illuminate (Moore 2017), so too can the most sophisticated models. In this paper we are considering the impact (and hence pitfalls) of relying on formal models – that is mathematical or computational models. This is what we will mean when we talk about "models" here; if we mean informal models, we will explicitly say so as in "mental models" or "informal models".

On balance, it may be argued that people are relatively good at reflecting and negotiating how we collectively think about things, via social and political processes - although it should be clarified here that we do not all necessarily possess the same negotiating capacities or opportunities to affect these processes in a similar manner. What is relatively new in these social and political negotiating arena(s), and what we need to carefully consider, is that formal models are increasingly being introduced into (and their legitimacy questioned within) these processes, in terms of their results or their underlying ideas. This paper looks at some of the dangers that this might introduce.

A complex simulation model can be a powerful tool – capturing and integrating knowledge that would be almost impossible to do in other ways, then facilitating calculations from that knowledge. However, these complex tools can be difficult to construct (adequately), and even more difficult to use (appropriately). Consequently, there is always an underlying chance of fooling yourself or others when building or using them, and hence the possibility of prompting bad decisions. Furthermore, complex models can act as a mistake amplifier, making small mistakes have big consequences. As the saying goes

> *"To err is human, but to really screw things up you need a computer."*

This paper looks at some of these pitfalls in the hope that we might raise awareness of them and their consequences. We have structured our account according to two phases – *firstly*, the construction phase (and all that goes or should go with that), *secondly*, the application phase when the model has been released to be used in the wider world.

In the event that models are to be used to inform policy, both of these stages merit close consideration, not only by those involved in developing the models, but also by those that will be affected by them. As such, we lay out some of the core issues that may arise within these stages, and highlight how these might present pitfalls for the modeller, for the policy maker, and for wider society. This rough categorisation is simply to aid the reader by giving them some structure and is not to be taken as definitive – any of the pitfalls might affect anyone.

Not everyone involved and implicated by this process is a modeller or policy maker. The shortcomings of the interaction between modellers, the modelling process, policy makers, and the entire policy process are likely to be felt most by many facets of society that may have had very little (or no) bearing on or input into this process. A wider awareness by stakeholders and the public of the pitfalls may encourage them to be more critical of model-informed outcomes and to direct debate more towards the options and decisions being made.

2 Constructing a model: Pitfalls for modellers to avoid and policy makers to ask about

There are a number of pitfalls that can occur in the construction phase of the model. Although there is no shortage of evidence of poor modelling practices, and the negative consequences these can entail for society, many aspects of models are not usually subject to close examination by people outside the original modelling team[1] (Saltelli & Funtowitz, 2014). As such, the points made here are aimed at modellers. However, they hint at the type of questions that those who are considering using a model (i.e. policy makers) should ask when presented with a model, regardless of how impressive it may look. Doing so may go some way to avoid *some* of the potential pitfalls outlined for the following phase of model application.

It is worth stressing, at this stage, what a model is. Models are abstractions, formal constructions that represent aspects of the world. They are created by someone, somewhere, for a particular purpose, most likely to answer a particular question, and most certainly drawing on various assumptions. Indeed, there's no getting away from the fact that in engaging in the activity of modelling assumptions about reality have to be made. While these assumptions might be more or less reliable (given the context and purpose of the model), there are potential traps that even the most experienced, competent, and respected modellers may fall into. The overarching point we wish to make is that a combination of reflexivity and transparency is key to avoiding some of these pitfalls.

2.1 Modelling Assumptions

[1] At best, examination is by a few in the same domain as themselves – people who likely have the same assumptions and world-view. Thus many models are not *effectively* critiqued in an independent manner.

How we choose to see the world is itself a complex and subtle process that is not well understood. How we view the world, how we choose to represent the world impacts and how we frame the world (Giampietro et al, 2013). How we conceive of a problem matters a great deal because it informs and limits how we try to solve that problem (Moore, 2017). The institutionalisation of seeing the world through numerical abstraction as the most authoritative way of seeing has a long history, and one that is deeply embroiled in facets of power (Bavington, 2009; 2011; 2015; Moore, 2017; Scott, 1998). It is especially interesting that post 2007 crash we continue to view words largely as 'interesting points of view', while numbers 'never lie' (or at least are more convincing), and difficult calculations or simulations often retain an unwavering authority. What is perhaps even more worrying in this, however, is that this institutionalisation runs so deep, and is so ingrained in our thought, that often we fail to recognise that numerical abstractions and complex models can be as laden with 'points of view' as other forms of knowledge. Failing to recognise this is the first pitfall that both modellers and policy makers are likely to fall into.

As per our aforementioned statement, all models are built on assumptions about reality, and this includes complex models. These assumptions are both implicit and explicit. Some may be based on theory, others on empirical evidence, maybe a mix of both, or, perhaps, something altogether more ad hoc, such as those derived from tradition. We may not even be aware of the implicit assumptions. These assumptions will determine what goes into a model and, perhaps more importantly, what we leave out of it (Sterman, 2002). Together these will have a bearing on our entire conceptualisation of the problem at hand.

Some of the assumptions we make will be somewhat reasonable, while others may be downright unreasonable. A good example of this comes from the models that are employed in fisheries management, some of which assume that nature is a stable system (although most current thinking acknowledges that it is anything but stable). The corollary of this is that goals or policies designed according to this assumption (e.g. Maximum Sustainable Yield (MSY), which has become the de-facto goal of most fisheries management regimes today) may turn out to be fairly unreasonable themselves (Bavington, 2015). Models populated with homogenous rational economic agents are another good example. Stiglitz (2011) has highlighted that many of the standard economic models so deeply implicated in the last financial crash had critical omissions, along with a raft of incorrect assumptions, over-simplifications, or the 'wrong' simplifications. In turn, policies that have been designed according to these models were, and continue to be, worryingly dysfunctional.

2.2 Theoretical spectacles

As indicated, many things can colour our perception of reality and thus the assumptions that we make – for a good discussion on this see Glynn (2017). Among the factors he highlights are experiential and environmental biases (including disciplinary biases) that – regardless of how 'objective' they view themselves – scientists are subject to. Their disciplinary orientation will most certainly entail some commitment to a worldview that leans towards some value systems over others or to depicting aspects of that worldview over others in their models. Thomas Kuhn described this effect as wearing "theoretical spectacles" (Kuhn 1962) – the theories one believes lead one to only notice the aspects of the world that fit the theories, and not those that do not.

This is often inadequately considered by those engaged in building such models. In fact, as Sterman (2002) has observed, narrow modelling assumptions are a common occurrence, even in work that has been published in highly respected journals. Although some complex models (e.g. agent-based simulations) allow the avoidance or widening of some of these assumptions, it would be wrong to think that such models are free from unconsidered or over-simplistic assumptions that may critically affect the results they give. Thus, a potential pitfall for modellers is failing to sufficiently consider and/or critique assumptions that underlie a model's construction and how useful, or dangerous, these may become down the line.

Modellers tend to spend a significant amount of time with their models – deeply engaged in constructing them, thinking about them, and tuning them. Thus the danger of "theoretical spectacles" is particularly acute for modellers as they often learn to see the world 'through' their models and begin fitting it to adhere to what they perceive to be true, developing a strong confirmation bias (Sterman, 2002). This process can result in modellers making models that fit 'their version of reality' quite well, but it may not necessarily reflect observed 'reality' very well[2]. What we perceive to be true is based on our assumptions, and if these are not subjected to sufficient independent or reflexive examination and critique, then dangerous, brittle or simply wrong assumptions may be included in the models we use. The upshot of this is that we get models built on bad foundations that may perform in a completely inadequate, indeed sometimes catastrophically mistaken, way.

2.3 Over-simplified models

For models to be understandable, they need to maintain a certain level of simplicity. However, it is always worth keeping in mind that while simplifying reality it is likely to become more removed from that which we are trying to represent. As such, it is helpful to be aware that a model might not necessarily prove encompassing enough to incorporate alternative understandings, experiences, values or needs of those whose reality is being abstracted. It merits consideration that in our abstractions, that which we exclude is likely to matter for something, or someone, somewhere. In this sense, it is worth remembering that highly stylised interpretations can work to colour our vision from the beginning (Moore, 2017) – the more stylised they are the stronger the 'colouring' might be (since it makes for a more attractive and portable story). Similarly, Glynn (2017) argues that highly simplified models may be poorly suited when confronting issues that we may not have experienced before or have little experience of.

Agent-based models are generally more straightforward in how they represent the world – allowing one computational entity for each actor, for example. This means that they do not need such simplifying assumptions as models, which represent populations as abstractions. However, they are still subject to the same pressures as other kinds of model, and there is a strong academic and publication bias towards simpler models.

Notwithstanding this, many models – even the simplest – can be useful, and one way or another we are all working off *some* kind of model (even if only a mental or informal model). Given that we cannot observe or measure everything, everywhere all of the time, simplifications can and do help us understand some complex processes that we may otherwise not understand (Glynn, 2017). The pitfall here for the modeller to avoid, however, is one of *over*simplification, and any type of model can fall guilty to this charge. Complex computational models can become oversimplified, if the modeller may be constrained by time limitations, information limitations, their aforementioned worldviews, computational capacities, and so on. The danger here arises when simplification leads to a level of abstraction that misses key mechanisms and aspects that really *are* important in the process we are trying to understand. For example, this might happen through choosing to restrict what goes into a model to available numerical data – because it is easier than dealing with non-numerical data (Sterman, 2002). Oversimplification, subsequently, can lead to many kinds of error, including: human errors, computer errors, incorrect/misleading results, biased or limited interpretations, and ultimately, to bad decisions (Glynn, 2017). Whilst it is very difficult to be sure as to what aspects are crucial to include, or to include every tiny nuance that might be relevant or important (even ethnographers struggle with this one), the task here is to be upfront about, and reflect on these simplifications, and try to catch them out if they are oversimplifications. At the very least simplifying assumptions should be clearly acknowledged and documented.

---

[2] This is shorthand for saying the model's assumptions make a model useless in terms of its purpose.

2.4 Underestimating model limitations

Given the two previous points that have been made – that models are built on assumptions about reality, and even the most complex entail a fair amount of simplification – it is not difficult to make the point that all models are going to have some limitations. The usefulness of a model is going to be constrained, so that an over optimistic selling of your model for any and all purposes is not likely to end well – for anyone. Building a representation of the human or natural world – or both – is hard, and it would be a mistake to think otherwise. Building an oversimplified/over prescriptive model and putting too much faith into what it can tell us can have many negative repercussions.

Some idea of a model's limitations can be gleaned through the assumptions that are built into it – it is unlikely to work well in situations where the assumptions do not hold. Other clues to a model's limitations can be found by running the model under many different considerations, e.g. its sensitivity analysis. However, the final arbiters of a model's limitations are usually only apparent when the model is used in practice. Thus, models need to be continually reviewed as to their continuing suitability and usefulness.

Particular care is needed when the model is being applied in a context that is very different to the one it was designed for or tested within. In a way, each time a model is applied in a different context its utility there should be separately established, and not taken for granted. The more different the situation the more it needs retesting, but this is often not done due to the cost of this. It is much easier just to reuse the model and hope for the best – easier in the short-run that is.

One subtle way that a model can be used beyond its limitations is when it is subsumed as a sub-model of a bigger, more complex model. In such cases, the failures of the model can be masked by all the other things going on and not noticed. However, since models can 'amplify' error and bias, it might have an even bigger impact on the results.

Thus it is important to remember that even the best models have limitations. Models are not (or 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.

2.5 Not checking and testing a model thoroughly

Clearly if there is any danger that a model might be used to inform real world decisions, then the modellers and/or model users have a duty to check and test the model for its intended purpose as carefully as possible.

3 Unleashing the model: pitfalls for modellers, policy actors and society.

Saltelli and Funtowitz (2014) have made the case that models should never make it onto the policy arena without undergoing rigorous and independent sensitivity auditing. This is a fairly reasonable suggestion, given that many of the pitfalls that can occur during the development phase (which should include testing) of the model can have fairly serious implications in the event that it is applied, and its results are taken seriously, with little question. There have been some very public examples of this over the past number of years, that have had implications for those who are in the business of developing models, those in the business of designing policy, and, in turn, those who have to live with the consequences of these policies (e.g. Cavero & Poinasamy, 2013; Cassidy, 2013; Pierce, 2008). One such example will be discussed in the final section. However, even in the event that best practice has been strictly adhered to prior to the model's application, there are still a number of pitfalls to be avoided at this stage of the process. Thus, some of the big traps to be navigated at this stage relate to the state of the evidence base, and confusion over what a model can deliver on and what it realistically cannot. Points which merit consideration here include the purpose for which the model was built, the conditions it was built to satisfy, high levels of uncertainty, and the inability to answer the less 'scientific' questions that are being asked with the required less 'scientific' answers.

3.1 From lack of evidential support, mistaken or misleading models, to entrenched but ill-informed policies.

As we have indicated, models are constructed using little snippets of information, about somebody, something, or some event or process that has been observed. It is worth reiterating here that even for the most complex models the 'real world' is going to be more complex than the pieces of information we have on it. This also holds true for 'good' models. However, if things go wrong in the construction phase you may well be dealing with a fairly inaccurate model. It is also worth reiterating here that sometimes models are not grounded in empirical evidence but rather in tradition, that is, disciplinary theories that may well never have been proved beyond theory. Some of these models may be relying not only on tacit, but wholly unverified assumptions (Saltelli & Funtowitz, 2014). It seems reasonable to suggest that models based loosely on real evidence, or perhaps none at all, are going to throw up problems at some stage if used in guiding the policy formation process. An easy target here are the aforementioned but often used, well-versed models on the policy scene coming from economics. While models such as these may make policy formation somewhat easier, in that they give a representation of something in a way that gives it an appearance of manageability, it is worth considering and questioning in detail their assumptions, and the *actual* wisdom that they encapsulate.

Unfortunately, history has taught us that mistaken or misleading models can quickly gain traction. Particularly, if they *appear* to offer a workable solution that is amenable to policy making (see the example in section 5.1 below). Furthermore, once these models (or the policies they have justified) become institutionalised – even when our knowledge has progressed so that we can see that the models (and thus the policies they inform) are underpinned by incorrect assumptions – it may be very tempting to continue to use them because the alternative is too messy or appears too hard. Essentially the pitfall here is that models can become so embedded within the policy making process that they are difficult to change. This may be for a variety of reasons, including that they reinforce particular interests or simply out of sheer habit. "We've been doing it this way for thirty years, so it must be right" is most certainly a pitfall policy makers can (and do) fall into (see Rosewell, 2017 p. 163). While this may make the game of policy making easier, it may not make for the best societal outcomes, and these may range from minor to fairly catastrophic consequences.

Fisheries management provides an example of this. Many of the models that have traditionally been employed within fisheries policy formation have become so deeply institutionalised, that they are still widely used within fisheries management (both in the EU, and elsewhere) – despite being increasingly questioned by scientists themselves, both in terms of the assumptions that underpin them and those that were left out (Bavington, 2009; 2011; 2015). This type of institutionalisation can work to weaken the evidence base, with deleterious effects. As Bavington (2015) argues in relation to the collapse of the Newfoundland cod, once a goal had been set according to these models, once it became the core objective of global fisheries management and institutionalised into international law, there was very little room left to question the suitability of these models. These models and their projections would, in turn, be implicated in the collapse itself (see section 5.2 below).

3.2 Model spread

One of the big advantages of formal models is that they can be copied and used extensively with little effort. This can have big advantages in terms of allowing others to inspect, critique and improve these models, but it also has downsides. One of these downsides is that models, once made and accepted in some way, tend to proliferate. That is, they tend to spread as if on their own accord. Of course, the ease of their reuse means that it is tempting to reuse them with little care or attention. In particular, care to retest or otherwise evaluate the applicability of a model for each area of application. In addition, once a model becomes widespread then others take this as a mark of its suitability, so that it spreads even more.

An example of this is when the "Black-Scholes" formula (Black & Scholes 1973) and its extensions became a common basis on which to price many kinds of financial derivatives. However, it later turned out that in other than the circumstances it was originally conceived for it gives misleading prices, e.g. in the presence of extreme price changes, long-term price variation or when dynamic hedging is not possible. The prevalence of this model has even been blamed for the bank crash of 2007/8 (Stewart 2016). As the famous investor, Warren Buffet, put it in a letter to shareholders "*The Black–Scholes formula has approached the status of holy writ in finance ... If the formula is applied to extended time periods, however, it can produce absurd results. In fairness, Black and Scholes almost certainly understood this point well. But their devoted followers may be ignoring whatever caveats the two men attached when they first unveiled the formula.*"[3].

3.2.1 Narrowing the base even further

Another way that models, once unleashed into the policy making process, can affect the evidence base is through narrowing it. The case of the Newfoundland cod, mentioned above, indicates how models can work to constrain the evidence base, therefore limiting decision making. In this sense, a policy maker pitfall would be narrowing the evidence base to the part, which is seen as authoritative, and all other evidence is side-lined. This raises further questions in relation to what a model may and may not be able to adequately capture, and whether there may be other sources of evidence better suited to that task, and the institutionalisation of what we deem to be authoritative evidence.

This point is very much related to our earlier point regarding the "theoretical spectacles" – we all wear some type of spectacles that have been coloured by our environment and our need to navigate it. A model might be built from one viewpoint using a particular set of scientific spectacles, and used in accordance with the different spectacles of a policy maker. These spectacles might bias or limit out vision in innumerable ways. While these limited viewpoints might be ok within their original context of development and use they may not adequately capture things outside of it. For example, this view might not be compatible with the spectacles those operating in the context, with which the policy is concerned or will be employed in, are wearing.

This, of course, remains a challenge in policy making today, despite the widespread rhetoric in favour of stakeholder engagement, participatory governance, and human dimensions. Pearce et al. (2014, p163) have made the case that the tendency to prioritise technical data (numerical data and the output of formal models) over all else is still a firm feature of the policy-making process. They highlight that studies indicate that the 'prevailing order' of the evidence-based policy process remains firmly rooted in traditional power hierarchies that are buttressed by a technocracy. In contrast, qualitative research and local knowledge are marginalised, so that a belief in the superiority of scientific methods from the natural sciences remains entrenched. Further, Saltelli and Giampietro (2017) have argued that modelling, when unleashed onto this space, can actually exacerbate this.

This becomes a pitfall in the sense that, even though models can help us to understand something in a way we previously were unable to, it might effectively limit consideration of other understandings 'out there' that are likely to require consideration, or perhaps might even trump the model itself.

[3] http://www.berkshirehathaway.com/letters/2008ltr.pdf (accessed 1st June 2017)

The danger then becomes that models may work to propagate established forms of thinking to the detriment of all others.

Arguably, given the special status we bestow on models (perhaps arising from their impressive appearance or the authority they gain from their scientific status) it is worth considering that these processes are imbricated in power in a number of ways. In this sense, the representations we present and use can work to cement this. The authoritative role of models can help justify the centralisation of decision making or perpetuate a top-down hierarchical mode of regulation. Modes of management, which many are increasingly recognising as suffering a crisis of legitimacy and from which they would like to move away. Given the special status this type of model can command, it can be used to justify decisions – decisions, which may not necessarily have the best outcomes.

Modellers, given the authoritative position of science, or at worst the appearance of science, can trip into the pitfall of further contributing and perpetuating these hierarchies, which may result in poorer answers than may have been available elsewhere. In this way, models may work to further exclude or obscure other ways of knowing; other ways that might prove to be a better answer to the current complex global challenges that policy makers, and society have to deal with. In this sense, they may perpetuate the failure to integrate or deal seriously with other forms of knowledge.

4 Some other things to be aware of

4.1 Confusion over model purpose

Good models will have, or should have, a clearly stated purpose – at least those that are applied to issues of real importance. Such a model will have been designed with that purpose in mind and tested with respect to this. If it is used for another purpose then it is likely to fail at this. Therefore, that model will only be able to help when used for its particular purpose, e.g. for scenarios where that kind of role is required. These kinds of confusion are dealt with in (Edmonds 2017).

4.2 Confusion over the kind of question a model can answer

A related confusion is when a model is designed to answer a question or inform thinking about one kind of issue is assumed helpful for a different question or issue. Take, for example, the bio-economic models of fisheries management. These models are built using biological and economic parameters, and largely ignore social parameters. They are designed with these objectives in mind. Proponents of these types of model are sometimes explicit about this, and may indicate that although other social objectives like employment, equality, or biodiversity conservation are important, they are not explicitly modelled (e.g. see Costello et al, 2016). As a policy maker, it is worth considering whether these models may be the most appropriate tool for suitable policy formulation, or to which the extent they should be relied upon. Interestingly, after years of managing fisheries based on these models, the poor social outcomes with respect to fisheries management are often downplayed or ignored, though, arguably, are not at all unexpected.

As a policy maker, one should be aware that a lack of clear purpose for a model is far too common (Edmonds 2017). It is therefore sensible to inquire carefully into this, and consider whether the purpose of the model is compatible with the kind of policy one is trying to design, and whether it meets one's objectives. Indeed, Sterman (2002) argues that, along with incorrect or missing assumptions, models often fail because more basic questions about the suitability of the model for the intended purpose were not asked. As such, the pertinent questions to ask become: Are the assumptions being made in line with the purpose of the model? What would this mean in terms of policy output? What kind of contradictions might this lead to?[4]

[4] Giampietro & Saltelli (2014) provide a discussion on these questions in relation to the Ecological Footprint.

A pitfall for modellers here would be failing to declare whether the model is suitable for the particular purpose, or whether it might hold up under different conditions, and falling into the trap of being too policy prescriptive on questions that are inherently political, rather than scientific. In this respect, some of the parameters/components we fail to include might matter a great deal, and might have longer-term societal consequences. Sometimes answering these questions will require much more than a model that has been built with specific, perhaps narrow, objectives in mind, using specific assumptions as to how society is, rather than how it could or ought to be.

4.3 When models are used out of the context they were designed for

Context matters! While a set of assumptions may accurately hold in one context, they might not in another – other factors could come into play in a new context that change the outcomes, or may even negate them entirely. For example, with bio-economic models in fisheries, scientists often acknowledge that the effects of their policy prescriptions, according to their models, assumptions, and goals, are likely to be context specific and depend on the social, economic, and ecological objectives within any given context. The danger here is that the policy maker is not adequately aware, or fails to consider this declared context-sensitivity, but rather goes off the tagline whereby the solution is posited without the necessary caveats. So although scientists may make certain caveats about their model explicitly clear, this does not necessarily mean that they are heard. Furthermore, these may be lost as the model moves up the chain to where policy will actually be implemented.

4.4 What models cannot reasonably do

It is worth highlighting that there are some things that models just cannot do. In these cases, the policy maker should not attempt to ask such questions of a model, and for a modeller not to present (spurious) answers if asked.

Some of the biggest questions we are trying to answer today simply cannot be answered by science, certainly not alone anyhow, no matter how much we dress them up with science (Weinberg, 1972). Three conditions give rise to such questions. Firstly, there are questions that science may not be capable of answering due to limited resources. Secondly, there are questions whereby the subject matter is just too variable to measure according to narrow positivistic frames (Weinberg explicitly places the social sciences as such a case). Thirdly, there are the types of questions or issues that involve moral and aesthetic judgements – they are not about 'facts' but values, although some questions may have elements of both (Weinberg 1972).

For example, if a policy maker asks a model to predict the consequences of a particular policy, and this is simply not predictable, then it is wrong to provide that prediction, even with caveats (because the modeller knows that the caveats will be ignored). If a policy maker tries to off-load the responsibility of a decision to the outcomes of a model, then this too should be resisted – it is the place of modellers to advise but policy makers to decide.

4.5 Uncertainty is too great

All models entail a level of uncertainty. This uncertainty usually increases exponentially with the complexity of the system we are trying to understand. A number of authors have highlighted this in relation to climate, and hence climate change. In this area, reasonable predictions are simply not feasible given the huge uncertainties this kind of modelling entails (Saltelli et al, 2015; Saltelli & Giampietro, 2017)[5]. Others have highlighted the total inaccuracy of these for comparing the possible damage (i.e. climate change costing) (Stern, 2016). There are simply too many processes involved here that we do not have an adequate understanding of, and as such, models of this kind ought not to be used for justifying policy decisions (Saltelli et al, 2015), and this is likely to stand regardless of

[5] However, model-based explanations of why climate change has been happening are well founded.

how super our 'super computers' get. The modeller pitfall that arises here is ignoring or hiding the uncertainties in their models, while for the policy maker it is allowing yourself to believe that we can quantify everything, including the uncertainty – which often we can't (Saltelli et al, 2015) – or failing to check that levels of uncertainty have been over or under estimated. Good science should be very cautious about even giving the impression that its outcomes are more accurate than they merit, but when the same scientists get involved in the policy process, there is a temptation to capitalise on their scientific status and use numbers or numerical representation to make their conclusions seem more certain and dramatic.

Saltelli and Funtowitz (2014) point out that a good indication that something may be suspicious about a model is if the information or numbers it offers up are too precise – something that provides the accuracy this implies just is not in line with what is usually possible with science.

Furthermore, engaging in this type of speculation has societal implications and throws up pitfalls for society, who may have little bearing on the actual policy process, by providing fuel for sceptics. Saltelli et al. (2015) show that introducing models into debates in relation to climate change may have done more harm than good, with the authors stating that society is potentially in danger of endless debates over uncertainties and competing arguments. The authors further highlight another, just as serious, issue – with excessive confidence in our ability to model the future, we may well commit to policies that reduce, rather than expand, available options and thus our ability to cope with what comes in the future.

4.6 A false sense of security

This point is interrelated with many of the previous points. As we have pointed out, just because models look very impressive or authoritative or provide us with a graspable number, they may not always actually be that impressive or authoritative, and the number may well be just as useful as one that was written down randomly on a sheet of paper. While at face value they might be quite enticing, they 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 the changes that are predicted. History has taught us that such an approach does not necessarily end well. Arguably, this is happening within fisheries, and many other domains as well.

Having a tool that can provide us with some kind of forecasts, there is the risk of relying on this, as a mechanism through which to avoid responsibility for perhaps the worst-case scenario or the unknown scenario. In this sense, we tend 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. For example, is it reasonable to assume that we can continue on our current growth trajectory, while solving the ecological crises we are facing and the increasingly social facets of each of these, and stay within the best-case scenario limit of, say, climate change?

4.7 Not more facts, but values!

Related to the points above there are just some things that a model cannot do, or compensate for. Models cannot provide us with or replace a moral or an ethical vision. They are unlikely, on their own, to provide us with a clear answer to the some of the hardest questions that the policy making process busies itself with (to avoid).

For example, a model may (somewhat) adequately capture some economic numbers, but it may not be able to capture what is important to real people. These personal values may be hard to capture or measure and so are not easily quantified. Examples of this may be cultural attachments to a place, which we might only garner through qualitative judgements. It is also unclear as to whether a model can really give a voice or include those that they seek to represent adequately – regardless of how participatory the approach employed has been. In this sense there are always going to be qualitative judgements to be made.

There is much scholarship based on looking at 'what could be', rather than drawing on models that look at 'what might be' based on assumptions about 'what is' – and this might be a more useful consultative tool for some issues. While these models may present us with some alternative course of action in relation to a specific question, they may not present us with any real alternatives for the future. If we are interested in articulating what 'could be' in a meaningful sense models may not prove to be very useful. So although a model may give us some sense of how things are, from a particular perspective, they might not be so good at telling us about how things should be, or how things could be (for an anthropological discussion related to this, see Holbraad et al., 2014).

5 Two Examples

5.1 An Economic Example

The 2008 crash and the recent financial crisis give ample evidence as to how models can go wrong – the pitfalls modellers can fall into, the pitfalls policy makers can fall (or jump) into when consulting models, and the severe societal consequences that this can entail. Particularly, when what turns out to be a seriously flawed model is used to justify particular policies with far-reaching and long-term consequences for the day-to-day lives of people. Of course, there are all sorts of modelling pitfalls that one might draw on in relation to the recent financial crisis, however, one really worth citing here, even though it has been widely cited elsewhere (e.g. Saltelli & Giampietro, 2015; 2017; Saltelli & Funtowitz, 2014) already, is the Reinhart & Rogoff case.

This case exemplifies what can go wrong – from dodgy assumptions and basic coding errors in the construction phase to the uptake (of flawed results) and implementation (along with institutionalisation) of very prescriptive policies. Policies that, once unleashed, resulted in really devastating societal outcomes. This example is particularly pertinent as it shows both the short term and longer term implications, many of which continue to reverberate in the lives of ordinary people today. It is at least partly imbricated in the current political climate, not only in Europe but on the other side of the Atlantic as well.

In 2010 Reinhart and Rogoff, two Harvard economists, published a study based on a model, that would provide the impetus for the implementation of severe austerity measures in many countries during the economic crisis. Their paper 'Growth in a Time of Debt' was widely publicised, and actively drawn on by policy makers. It argued that high debt had a negative impact on growth, and once this passed a threshold of 90% this would, potentially, become dangerous and actually impede growth (Cassidy, 2013; Reinhart & Rogoff, 2010; Saltelli & Funtowitz, 2014; Saltelli & Giampietro, 2015; Saltelli & Giampietro, 2017).

This was taken up by debt-facing policy makers on both sides of the Atlantic, and subsequently used to justify huge cuts in government spending and the implementation of austerity measures and packages in some countries, particularly across the EU (Cassidy 2013). A 2013 report by Oxfam highlights the shift towards austerity that occurred in 2010, which marked a turn from earlier more interventionist approaches to the crisis. In the UK, for example, prior to 2010 the government had taken the track of implementing a stimulus package, which included increased spending on social housing and education. This contrasted with post 2010 spending cuts aimed at reducing the deficit. Cassidy (2013) details references to the Reinhart & Rogoff paper being made by George Osborne in the House of Commons. Similar changes to public spending were implemented across the Eurozone, and elsewhere.

Three years later the work of Reinhart and Rogoff was replicated, and it turned out to contain some basic errors (Cassidy, 2013). The authors of this replication, Herndon, Ash and Pollen (2013) found: '*that selective exclusion of available data, coding errors and inappropriate weighting of summary* *statistics*' had led to serious miscalculations and inaccurate representations with respect to the relationship between public debt and growth. Unfortunately, this came too late for the people who were subjected (and continue to be subjected) to the policies the original paper had justified (Cassidy, 2013; Saltelli & Funtowitz, 2014; Saltelli & Giampietro, 2015; Saltelli & Giampietro; 2017). Indeed, many countries within the EU, under a great deal of 'encouragement' from the EU, have now institutionalised austerity via changes to legal mechanisms that mandate a balanced budget (Bruff 2016).

Austerity policies have had some wide-ranging effects. A 2013 Oxfam report (Cavero & Poinasamy 2013) documented the implications of the austerity programmes that have been implemented across Europe, arguing that with inequality and poverty on the rise Europe is facing a lost decade, with an additional 15 to 25 million people facing the prospect of living in poverty by 2025 if austerity measures continue. Indeed, decreased provision of public services, regressive taxation policies, rising inequality, persisting unemployment (in particular youth unemployment in some countries), increased food insecurity (as seen in the widespread appearance of food banks), health implications, lower income, debt burdens, and widespread discontent (Bruff, 2016). Perhaps some of the worst of these effects have been felt within the 'bail out' countries like Greece, but they have certainly not been restricted to these countries. These effects have been felt, and continue to be felt at the individual, household, societal, and wider political level, even within those countries now drawn as 'good examples' of austerity, such as Ireland (Bruff 2016). In this light, the polarisation of politics that we have increasingly seen no longer looks surprising. While there are obviously other things at play here (and we can not just blame the model), the model certainly is in some way culpable for how this has played out.

This story does not prove that the opposite policies would have turned out differently. We also do not know that the politicians involved would not have pursued the same policies anyway. What it does show is how sloppy modelling can be used to justify the policies that politicians choose, giving these more credibility than they might otherwise have, help to insulate them against criticism and debate, and thus to institutionalise the choices.

5.2 A Socio-Ecological Example

Bad models are, of course, not only confined to the world of economics and finance (although at times it may seem that way). The second example we draw on is one from fisheries management, with respect to the collapse of the Newfoundland cod. This example, again, gives us a sense of many of the issues and how they overlap. It also serves to illuminate how models can serve to override other sources of knowledge.

The story of fisheries and their often lamented status is of course a straightforward one – overfishing, or 'too many fishermen catching too much fish'. This narrative, while often cited, is problematic, in that it gives us little insight into the reality of modern fishing, and all that goes with it, including the way modern fisheries are managed. Certainly, it gives us little indication of the historical, political, economic, and ecological contexts of these endeavours and the relations that structure them. Within this world of fish, fishermen, fisherwomen, scientists, and managers (state and increasingly non-state actors) models feature highly – from population models of fish stocks and bio-economic models of efficiency to increasingly complex models of a variety of aspects of fisheries, including agent based models.

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 ended. The cod were declared commercially extinct (Bavington, 2011). What happened?

Much work has been done in this area, and many predictable answers have been put forth. A lot of these tell a simple story of overfishing, environmental conditions, and poor management. However, a number of authors (Bavington, 2010; Finlayson, 1994) have looked at the role of fisheries science in this story, arguing it and its models played a pivotal role in the collapse of the Newfoundland cod stocks. As the fishery and the fisheries management surrounding it developed it became about counting how many fish there were in the sea (Finlayson, 1998), and predicting how many fish could be caught, which was fed back to the scientists engaged in making these predictions. This, in turn, resulted in the development of increasingly intricate mathematical models based on differential equations. Partly due to the traditions in the field and the increasingly complex data they were trying to fit, these became more and more divorced from reality during this development (Finley, 2012, in Bavington, 2015).

Finlayson (1994) details the series of scientific blunders (based on models) that were made in the years leading up to the moratorium, in spite of repeated concerns being raised by inshore fishermen, with respect to the status of the cod. Despite a number of Commissions, and corresponding reports investigating the status of the cod in the years leading up to the collapse, despite protestations from the inshore sector they were seeing declining catches, the science depicted an increasing resource base. Successive failures were made in making adequate inferences in relation to the overall stock health – for example the Kirby Report indication that any reported decrease in profitability was merely down to cost-price squeeze. This report led to more development and investment in the fishery driven both by the state and individuals. Scientists and the fisheries department throughout much of the 1980s estimated a 15% annual rate of growth in the stock – figures that were consistently slated by inshore fishermen. Similarly, the subsequent Alverson Commission was formed to investigate the declines being reported by inshore fishermen, but cited environmental influences on the annual inshore migrations of the stock. Again such findings were contested by the inshore sector.

It was not until 1989 that this erroneous forecast for fish stocks was corrected. The fisheries department issued its annual assessment based upon revised mathematical models to generate stock estimates from research and catch data, which indicated that abundance had been overestimated by as much as a factor of two. The subsequent Harris Commission found that the fisheries department's estimates of stock strength were based upon data, methodologies, and models of such poor or uncertain quality as to be essentially useless as a rational basis for management or commercial planning.

The executive summary of the Harris Report (1999:2) states that:

> "*During the next seven years the euphoria that had been engendered by the declaration of the exclusive economic zone was reinforced by the steady growth of the stock, by continually improving catches, and by the belief that the FO.I objective was, indeed, being met. In those circumstances, 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 (1994, pp. 12-15) argues that social dynamics were certainly at play in generating some of the stock assessments in this case. This points out that in this instance, scientists and policy makers had become so committed to their description of reality (despite its wild inaccuracy), and: " *the idea of a strongly rebuilding Northern cod stock that was so powerful that it can be shown to have been read back into ambiguous data through analytical models built upon necessary but hypothetical assumptions about population and ecosystem dynamics. Further, those models required considerable subjective judgement as to the choice of weighting of the input variables*" (Finlayson 1994, p.13).

6 Conclusion

There are many pitfalls in both the modelling and policy arenas, and many of these feedback upon one another. However, each of these arenas has its own experts and professionals who will (hopefully) be aware of their own kinds of pitfall. It is perhaps when the policy and modelling world interact that many of the worst mistakes are made: when the policy actors do not understand the models or when the modellers do not understand, or assume adequate responsibility for, the consequences of their modelling. Thus, particular care needs to be taken when describing the capabilities or reliability of models to non-modellers, and policy actors need not to delegate their decision-making to a complex model that they do not understand, but retain their critical faculties.

It is also worth highlighting that the efficacy of a model is likely to depend on the question under investigation. Technical questions may not pose such a problem, however more complex problems will likely increase the urgency of the points that have been raised above, and are likely to require more information and wider consideration than simply drawing on *a* model.

The demarcation line between these two worlds is blurry in more ways than most of us like to admit, and this is not something new. As Weinberg (1972) pointed out:

> "*The politician, or some other representative of society, is then expected to say whether the society ought to proceed in one direction or another. The scientist and science provide the means; the politician and politics decide the ends – this view of science is of course oversimplified. Ends and means are hardly separable no matter how straight forward the question.*" (Weinberg, 1972).

However, the interface between scientists/modellers and the policy world is one that has increasingly come under scrutiny, and is one with a seemingly ever-increasing permeability. We suggest the following for complex systems where the science is immature:

- Stop using the word predict and stop expecting the word predict. Be very sceptical about any models that claim to be able to predict more than anything else.
- Use models to increase the number of alternative futures that might occur, rather than to reduce the apparent uncertainty.
- Ensure that models are re-evaluated frequently, especially when being used in a new context.
- Make effort to ensure that the models, the assumptions they are made from, and the whole policy process are open to scrutiny from all those affected.
- Even when a model is helpful by informing the formulation of a good policy, it cannot *decide* the policy. Deciding a policy is, and should remain, a political and not a technical process.
- Try and ensure that research and models that focus on what is happening now do not distract from the question of what 'could be' – the choices we have for the future.

The worst-case scenario is: What if these models or the ways they have been taken up are just *plain wrong*? Society has to live with the consequences. Thus, it is also important to remember we are not using these models in a vacuum; we are using them in a social, economic, political and environmental context that involves complex relations and power hierarchies. Ignoring the context and just focussing on the technical aspects of modelling may lead to bad outcomes for everyone.

References

Bavington, D. (2009). Managing to endanger: creating manageable cod fisheries in Newfoundland & Labrador, Canada. *Maritime Studies (MAST),* 7, 99-119.

Bavington, D. (2011). *Managed annihilation: an unnatural history of the Newfoundland cod collapse.* UBC press.

Bavington, D. (2015). Marine and freshwater fisheries in Canada: uncertainties, conflicts, and hope on the water*,* in Mitchell, B (Ed.) *Resource and environmental management in Canada*. OUP, Canada.

Benessia. A, Funtowicz. S, Giampietro. M, Guimarães Pereira. A, Ravetz. J, Saltelli. A, Strand. R and van der Sluijs. J. (2016). *Science on the Verge.* Tempe, AZ: The Consortium for Science, Policy and Outcomes at Arizona State University.

Black, F. & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*. 81, 3, 637–654.

Bruff, I. (2016) Scandalous Obfuscations of crisis in Europe, in *Progress in Political Economy (PPE),* available at http://ppesydney.net/scandalous-obfuscations-of-crisis-in-europe/

Cavero, T., & Poinasamy, K. (2013). *A cautionary tale: The true cost of austerity and inequality in Europe*. Oxfam International.

Christensen, V., & Walters, C. J. (2004). Ecopath with Ecosim: methods, capabilities and limitations. *Ecological modelling,* 172*(2), 109-139.*

Costello, C., Ovando, D., Clavelle, T., Strauss, C. K., Hilborn, R., Melnychuk, M. C., Branch, T.A., Gaines, S.D.,Szuwalski, C.S, Cabral, R.B, & Rader, D. N. (2016). Global fishery prospects under contrasting management regimes. *Proceedings of the National Academy of Sciences*, 201520420.

Edmonds, B. (2017): Five different modelling purposes. In Edmonds, B. & Meyer, R. (eds.) Simulating Social Complexity – a handbook, 2nd Edition. Springer.

Finlayson, A. C. (1994). *Fishing for truth: A sociological analysis of northern cod stock assessments from 1977 to 1990*. St. John's, Newfoundland: Institute of social and economic research, Memorial University of Newfoundland.

Funtowicz, S. O., & Ravetz, J. R. (1993). Science for the post-normal age. *Futures*, *25*(7), 739-755.

Giampietro, M., Allen, T. F., & Mayumi, K. (2006). The epistemological predicament associated with purposive quantitative analysis. *Ecological Complexity*, *3*(4), 307-327.

Giampietro, M., & Saltelli, A. (2014). Footprints to nowhere. *Ecological Indicators*, *46*, 610-621.

Glynn, P. D. (2017). Integrated Environmental Modelling: human decisions, human challenges. *Geological Society, London, Special Publications*, *408*(1), 161-182.

Harris, L (1990) *Independent Review of the Northern Cod Stock: Executive Summary, and Recommendations*, available at http://www.dfo-mpo.gc.ca/Library/114277.pdf

Herndon, T., Ash, M., & Pollin, R. (2014). Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge journal of economics*, *38*(2), 257-279.

Holbraad, M., Pedersen, M. A., & de Castro, E. V. (2014). The politics of ontology: Anthropological positions. *Cultural Anthropology*, *13*.

Karanikolos, M., Mladovsky, P., Cylus, J., Thomson, S., Basu, S., Stuckler, D., Mackenbach, J.P., & McKee, M. (2013). Financial crisis, austerity, and health in Europe. *The Lancet*, *381*(9874), 1323-1331.

Kentikelenis, A., Karanikolos, M., Reeves, A., McKee, M., & Stuckler, D. (2014). Greece's health crisis: from austerity to denialism. *The Lancet*, *383*(9918), 748-753.

Moore, J. W. (2017). The Capitalocene, Part I: On the nature and origins of our ecological crisis. *The Journal of Peasant Studies*, *44*(3), 594-630.

Pearce, W., Wesselink, A., & Colebatch, H. (2014). Evidence and meaning in policy making. *Evidence & Policy*, *10*(2), 161-165.

Rogoff, K., & Reinhart, C. (2010). Growth in a Time of Debt. *American Economic Review*, *100*(2), 573-8.

Rosewell, B. (2017). Complexity Science and the Art of Policy Making, In Johnson, J., Ormerod, P., Rosewell, B., Nowak, A., & Zhang, Y. C (Eds.). *Non-Equilibrium Social Science and Policy* (pp. 159-178). Springer International Publishing.

Saltelli, A., & Funtowicz, S. (2014). When all models are wrong. *Issues in Science and Technology*, *30*(2), 79-85.

Saltelli, A., Stark, P. B., Becker, W., & Stano, P. (2015). Climate Models as Economic Guides Scientific Challenge or Quixotic Quest?. *Issues in Science and Technology*, *31*(3), 79-84.

Saltelli, A., & Giampietro, M. (2017). What is wrong with evidence based policy, and how can it be improved?. *Futures*

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.

Sterman, J. D. (2002). All models are wrong: reflections on becoming a systems scientist. *System Dynamics Review*, *18*(4), 501-531.

Stern, N. (2016). Economics: Current climate models are grossly misleading. *Nature*, *530*(7591), 407-409.

Stewart, I. (2012) The mathematical equation that caused the banks to crash. The Observer, February 12, 2012, https://www.theguardian.com/science/2012/feb/12/black-scholes-equation-credit-crunch (accessed 1st June 2017)

Stiglitz, J. E. (2011). Rethinking macroeconomics: What failed, and how to repair it. *Journal of the European Economic Association*, *9*(4), 591-645.

Weinberg, A. M. (1972). Science and trans-science. *Minerva*, *10*(2), 209-222.