The Pitfalls of ABM

depending on your model purpose

Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan University

Introduction

Exploratory vs. **Justification Phases**

- It is normal (and useful and fun) to explore simulation models – that is, play around with them to get a feel for the kinds of behaviour that might result from different mechanisms and structures
- But this should be kept separate from when you get '*serious*' and want to use a simulation to justify a claim or argument that you make to others
- Then, in order not to waste their time, you need to be as clear as you can about about everything, including: aims, evidence, code, runs etc. etc.
- This is part of being scientifically rigourous

Modelling Purpose

- One crucial aspect is what *kind* of claim you are making using the simulation – what I call the *modelling purpose*
- This frames all the modelling work – since in public what you need to do is:
  1. Make your claim completely clear
  2. Use the simulation to support this claim
- Due to its fundamental role, this will effect how you build, check, run, document, and present your simulation
- Much confusion and *bad science* comes down to not being clear about this

Identifying then Mitigating for Potential Errors and Weaknesses

- Different kinds of modelling project (or purpose) can go wrong in different ways
- The approach suggested here is:
  1. Consider the 'threats' – the things might go wrong in pursuing that purpose
  2. Test and mitigate for these threats
  3. Report clearly on the threats and the extent to which you have ruled them out or mitigated for them
- Modelling complex social phenomena is very difficult – to make progress we have to be much more honest and careful about claims made

Modelling Purposes Covered

There are lots of different possible reasons to do simulation modelling (see Epstein 2008 in JASSS for 17 of them), but here we will only consider:

1. Prediction
2. Explanation
3. Theoretical Exploration
4. Illustration
5. Analogy

For each I define them, give examples, talk about threats and possible mitigating measures

Purpose 1: **Prediction**

Motivation

- If you can *reliably* predict something about the world (that you did not already know), this is undeniably useful…
- ...even if you do not know *why* your model predicts (e.g. a black-box model)!
- But it has also become the 'gold standard' of science…
- ...becuase (unlike many of the other purposes) it is difficult to fudge or fool yourself about – if its wrong this is obvious.

Predictive modelling

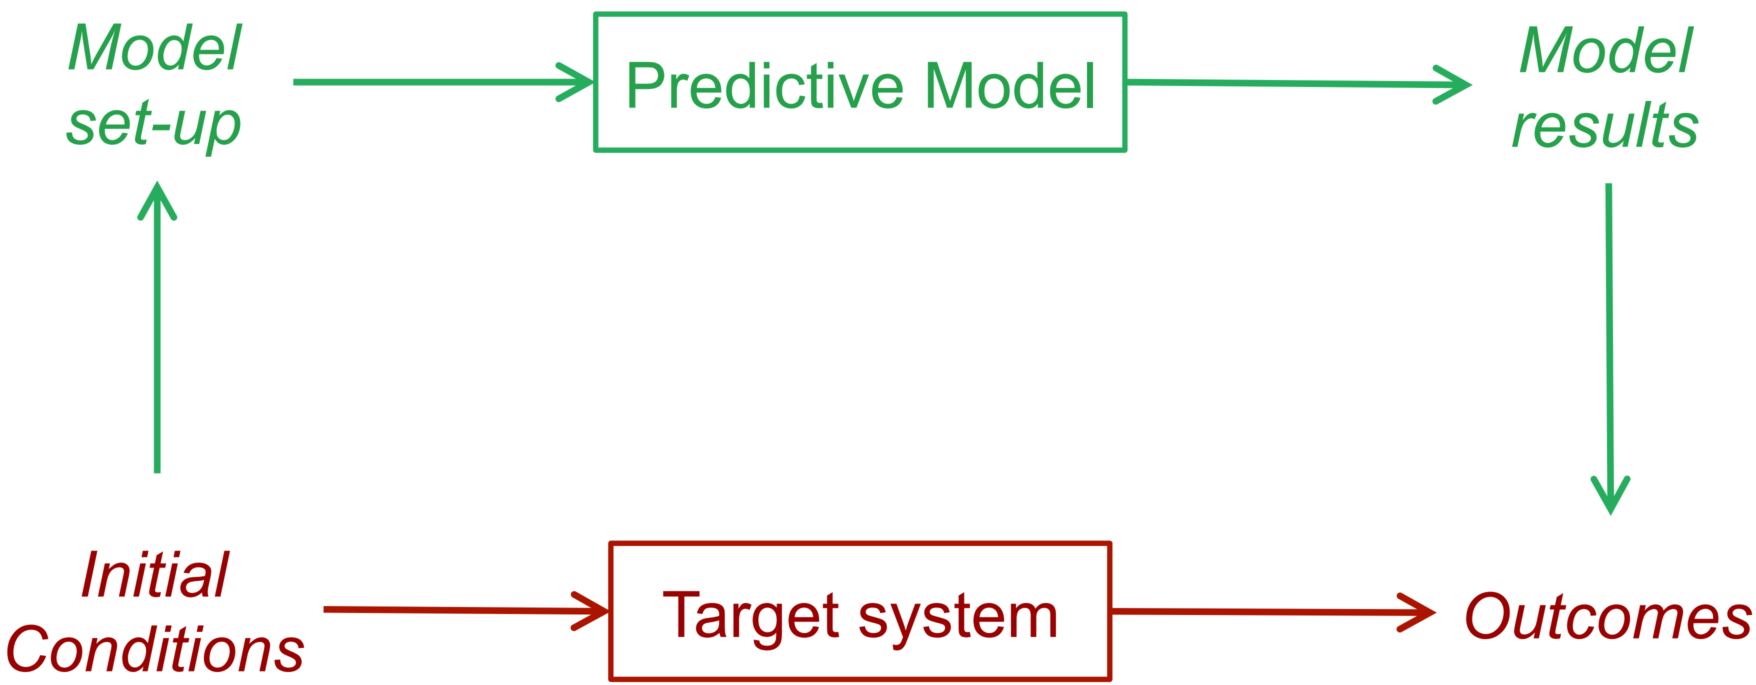

(Hesse 1963)

What it is

The ability to anticipate unknown data reliably and to a useful degree of accuracy

- Some idea of the conditions in which it does this have to be understood (even if this is vague)
- The data it anticipates has to be unknown to the modeller when using the model
- What is a useful degree of accuracy depends on the purpose for predicting
- What is predicted can be: categorical, probability distributions, ranges, negative predictions, etc.

Examples

- The gas laws (temperature is proportional to pressure at the same volume etc.) predict future measurements on a gas without any indication of why this works
- Nate Silver's team tries to predict the outcome of sports events and elections using computational models. These are usually probabilistic predictions and the predicted distribution of predictions is displayed (http://fivethirtyeight.com and Silver 2013)

Risks and Warnings

- There are two different uses of the word ‘predict’: one as above and one to indicate *any calculation* made using a model (the second confuses others)
- This requires *repeated* attempts at anticipating *unknown* data (and learning from this)
- because it is otherwise impossible to avoid ‘fitting’ known data (due to publication bias etc.)
- If the outcome is unknown and can be unambiguously checked it could be predictive
- Prediction is ***VERY*** hard in the social sciences – for this reason, it is rarely done

Mitigating Measures

- The following are documented:
  - what aspects it predicts
  - roughly when it predicts well
  - what degree of accuracy it predicts with
- Check that the model predicts on several independent cases
- Ensure the program is distributed so others can independently check its predictions

Purpose 2: **Explanation**

Motivation

- When one wants to understand *why* or *how* something observed happens
- One makes a simulation with the mechanisms one wants and then shows that the results fit the observed data
- The intricate workings of the simulation runs support an explanation of the outcomes in terms of those mechanisms
- The explanation is usually an abstraction of the model workings, so as to be comprehensible to us (e.g. a hypothesis about model behaviour)

What it is

Establishing a possible causal chain from a set-up to its consequences in terms of the mechanisms of a simulation

- The causation can be deterministic, possibilistic or probabilistic
- The nature of the set-up constrains the terms that the explanation is expressed in
- Only *some* aspects of the results will be relevant to be matched to data
- But how the model maps to data/evidence is explicitly specified

Explanatory modelling

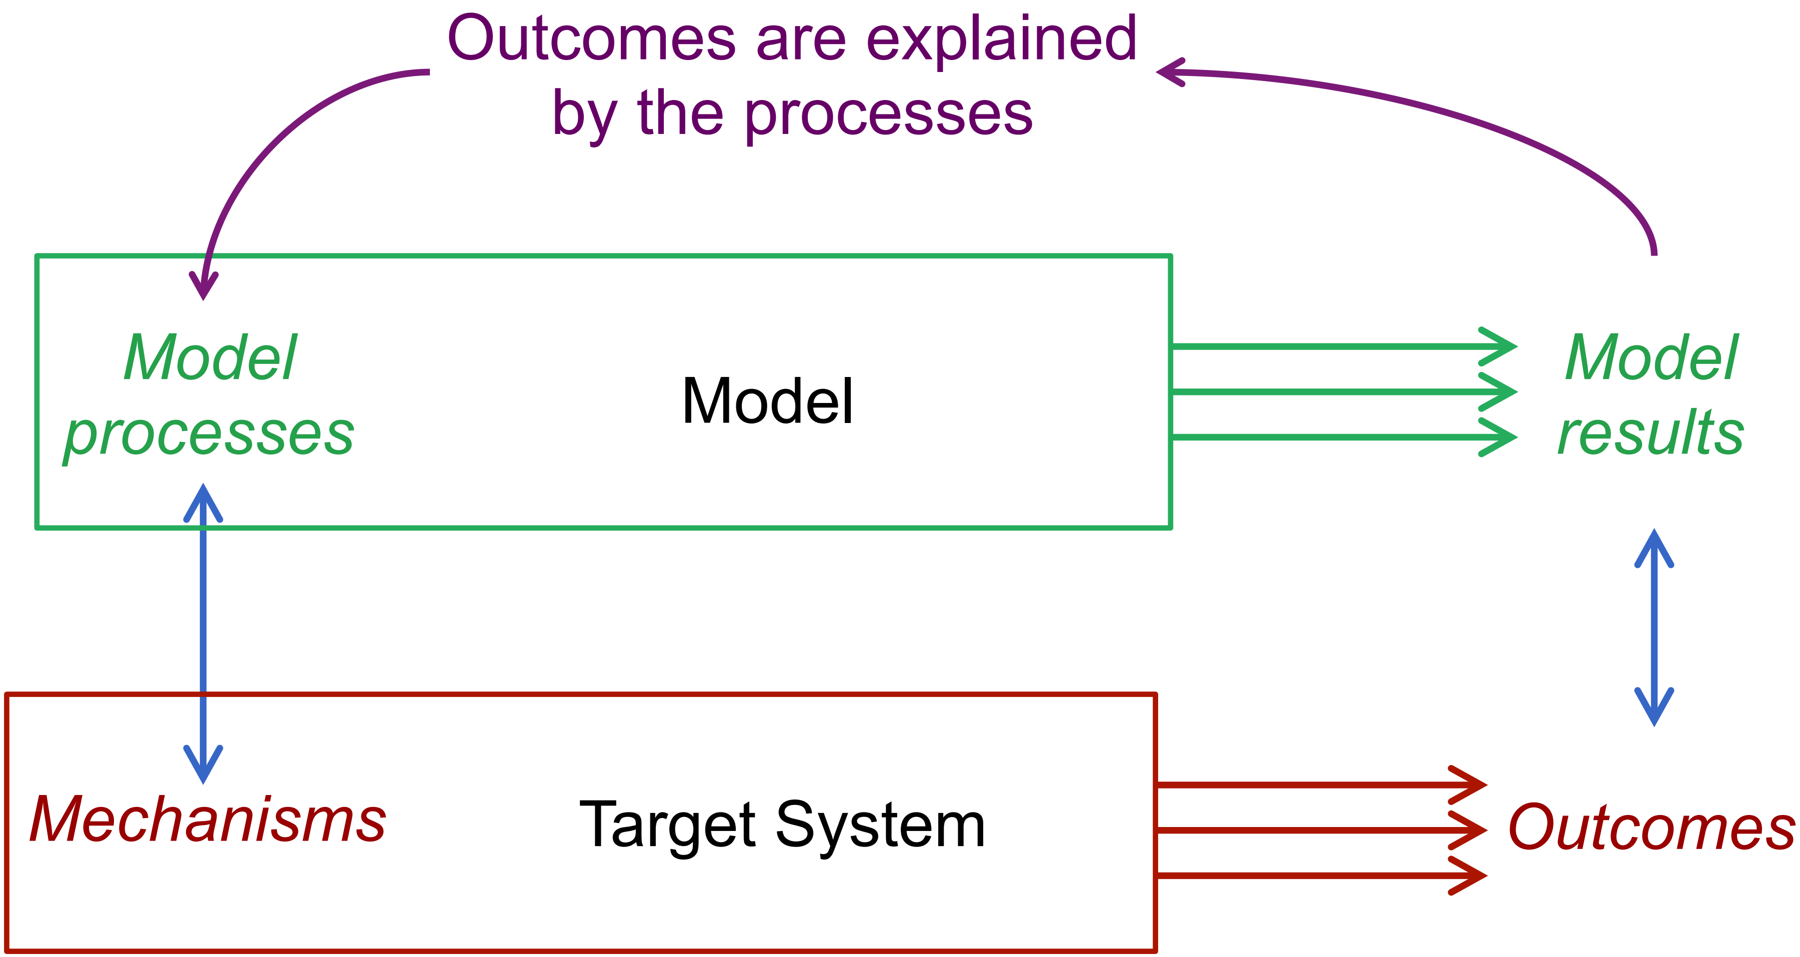

Examples

- The model of a gas with atoms randomly bumping around explains what happens in a gas (but does not directly predict the values)
- Lansing & Kramer's (1993) model of water distribution in Bali, explained how the system of water temples act to help enforce social norms and facilitate a complicated series of negotiations

Risks and Warnings

- A bug in the code is fatal to this purpose if this could change the outcomes substantially
- The fit to the target data maybe a very special case which would limit the likelihood of the explanation over similar cases
- The process from mechanisms to outcomes might be complex and poorly understood. The explanation should be clearly stated and tested. Assumptions behind this must be tested.
- There might well be more than one possible explanation (and/or model)!

Mitigating Measures

- Ensure the built-in mechanisms are plausible and at the right kind to support an explanation
- Be clear which aspects of the output are considered significant (i.e. those that are explained) and which artifacts of the simulation
- Probe the simulation to find when the explanation works (noise, assumptions etc)
- Do classical experiments to show your explanation works for your code

Purpose 3: **Theory Exposition**

Motivation

- If one has a system of equations, sometimes one can analytically solve the equations to get a general solution (i.e. a ‘closed form’ solution)
- When this is not possible (the case for *almost all* complicated systems) we can calculate specific examples – i.e. we simulate it!
- Using multiple runs, we aim to sufficiently explore the whole space of behaviour to understand the effect of this particular set of abstract mechanisms
- We might approximate these with equations (or a simpler model) to check this understanding

What it is

Discovering then establishing (or refuting) hypotheses about the *general behaviour* of a set of mechanisms

- The hypotheses may need to be discovered
- But, crucially, showing the hypotheses hold (or are refuted) by the set of experiments
- There needs to be a wide (maybe even complete) exploration of outcomes
- The hypotheses need to be quite general for the exercise to be useful to others
- Does not say anything about the observed world!

Modelling to understand Theory

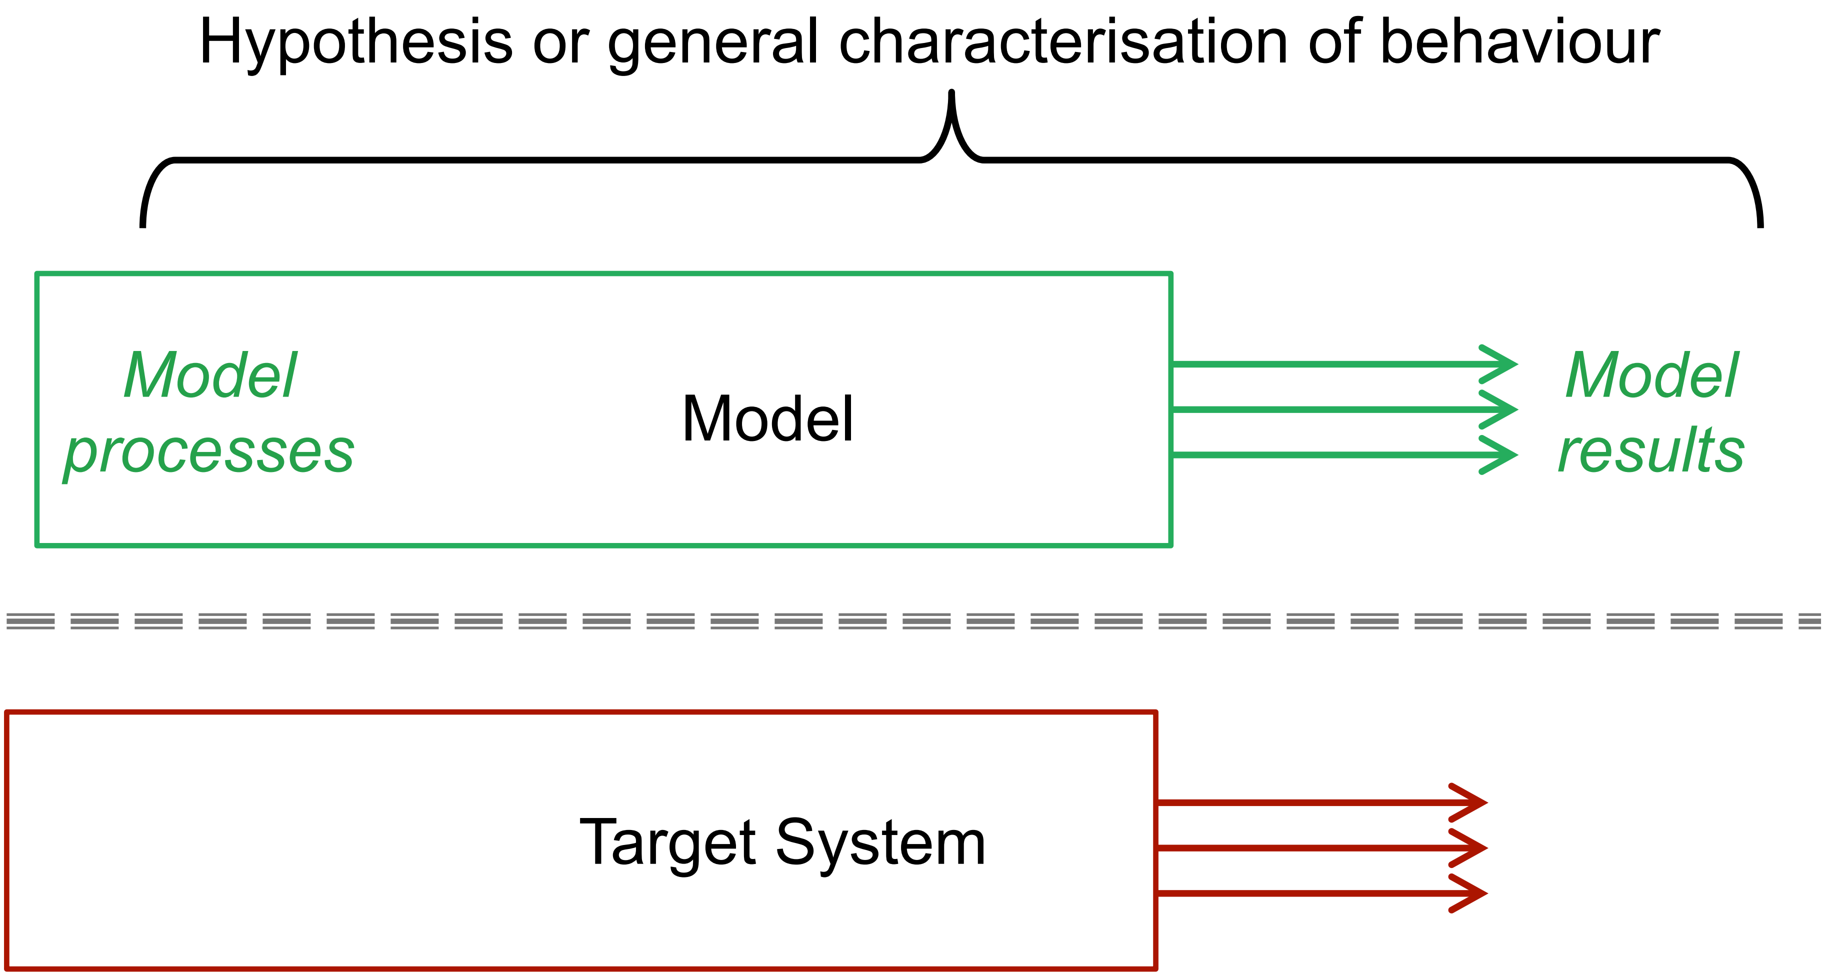

Examples

- Many economic models are explorations of sets of abstract mechanisms
- Deffuant, G., et al. (2002) How can extremism prevail? jasss.soc.surrey.ac.uk/5/4/1.html
- Edmonds & Hales (2003) Replication… jasss.soc.surrey.ac.uk/6/4/11.html

Risks, Warnings & Mitigation

- A bug in the code is dangerous to this purpose since the explanation might partly be based on an understanding of what the code was to do
- A general idea of the outcome behaviour is needed so the exploration needs to be extensive
- The code needs to be available so that people can test its assumptions etc.
- Clarity about what is claimed, the model description etc. is very important

Purpose 4: **Illustration**

Motivation & What it is

- An idea is new but has complex ramifications and one wants to simply illustrate it
- This is a way of communicating through a single (but maybe complex) example

A behaviour or system is illustrated precisely using a simulation in an understandable way

- It might be a very special case, no generality is established or claimed
- It might be used as a counter-example
- Simpler models are easier to communicate and hence often make more vivid illustrations

Examples

- Sakoda/Schelling's 2D Model of segregation which showed that a high level of racial intollerance was not necessary to explain patterns of segregation
- Riolo et al. (2001) Evolution of cooperation without reciprocity, Nature 414:441-443.
- Baum, E. (1996) Toward a model of mind as a laissez-faire economy of idiots.

Purpose 5: **Analogy**

Motivation & What it is

- Provides a 'way of thinking about' stuff
- The model does not (directly) tell us about anything observed, but is *about* ideas (which, in turn, may or may not relate to something observed)
- It can suggest new insights – e.g. new hypotheses or future research directions
- We need analogies to help us think about what to do (e.g. what and how to model)
- They are unavoidable
- They are useful, but can also be very deceptive

Intuitive understanding expressed in normal language

Common-Sense Comparison

Observations of the natural system of concern

Intuitive understanding expressed in normal language

Models of the processes in the system

Empirical Comparisons

Data obtained by measuring the system

Observations of the natural system of concern

Intuitive understanding expressed in normal language

Common-Sense Comparison

Models of the processes in the system

Observations of the natural system of concern

Examples

- Axelrod's Evolution of Cooperation models (1984 etc.)
- Hammond & Axelrod (2006) The Evolution of Ethnocentrism. Journal of Conflict Research
- Many economic models which show an 'efficient' market
- Many ecological models showing how systems reach an equilibrium

Warnings

- When one has played with a model the whole world looks like that model (especially to the model builder)
- But this does not make this true!
- Such models can be very influential but (as with the economic models of risk about lending) can be very misleading
- *At best*, they can suggest hypotheses about the observed world, but they don't demonstrate *anything*

Conclusion

How to decide what a model purpose is

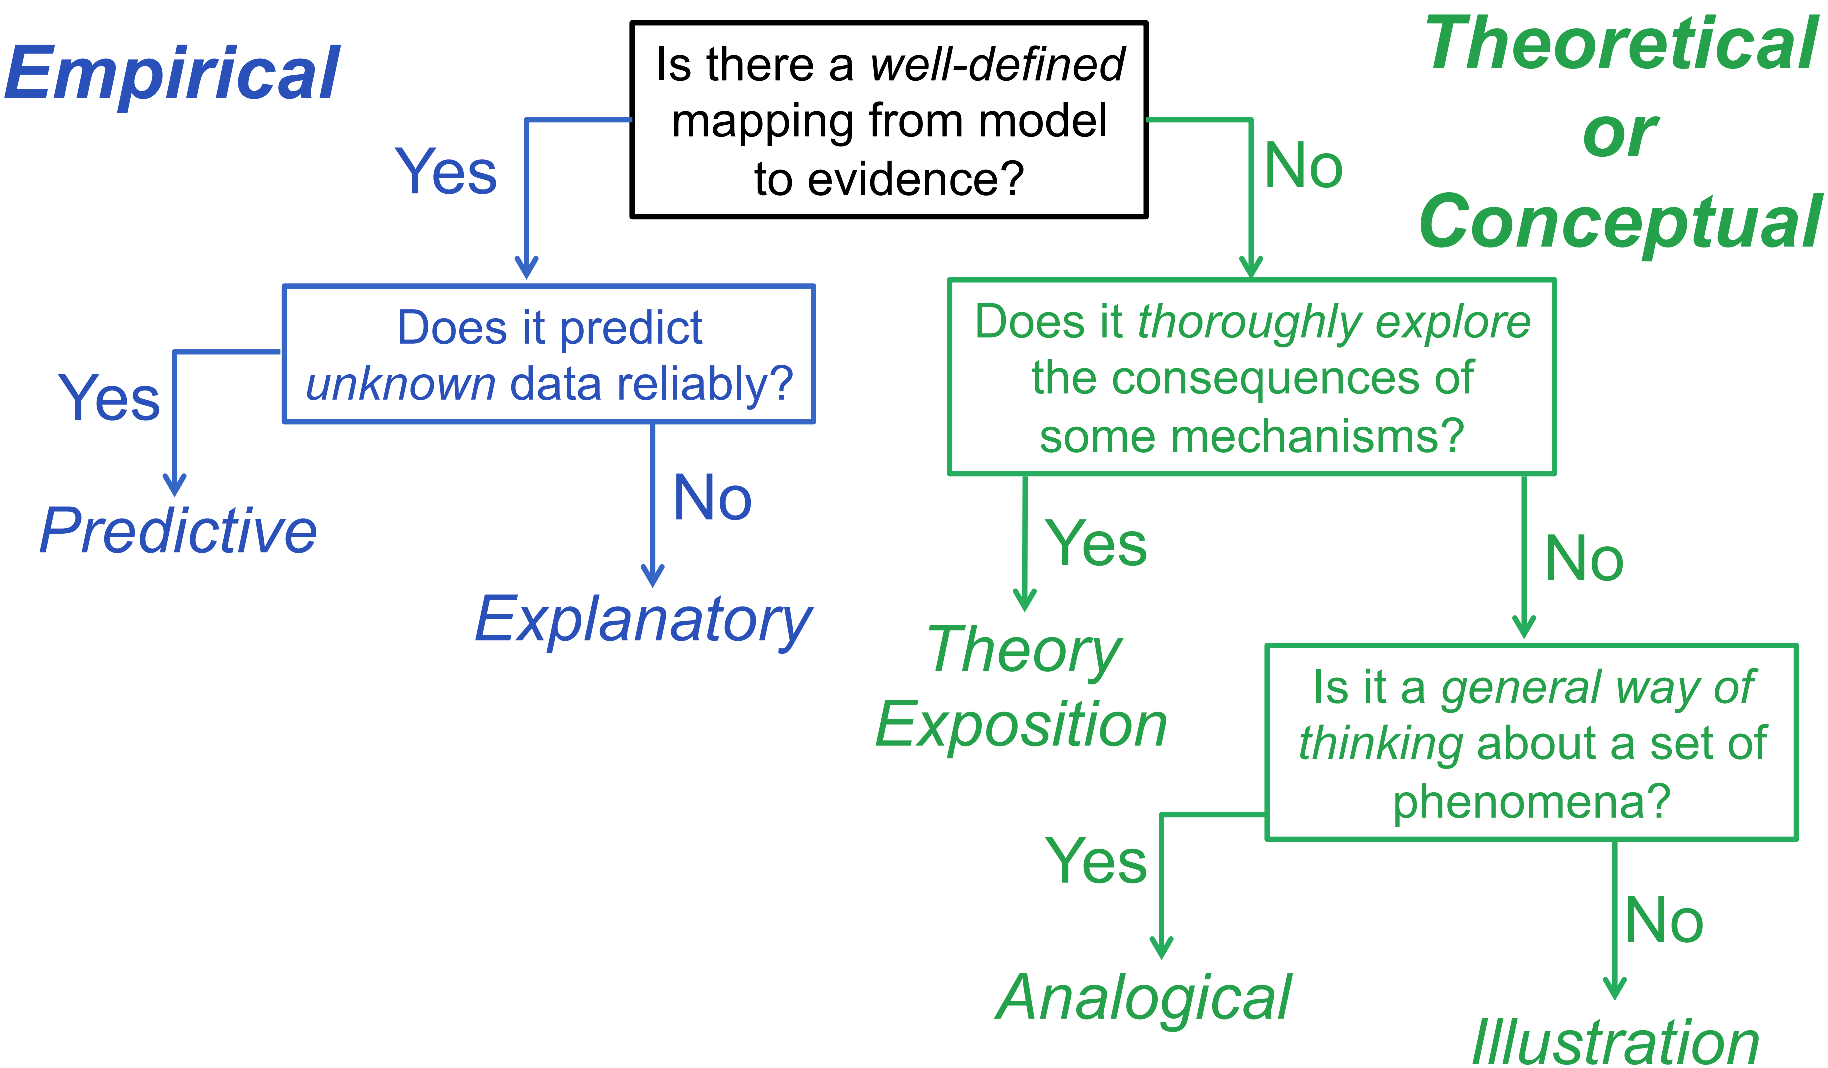

Summary of Purposes, features and risks

Modelling Purpose	Essential features	Particular risks (apart from that of lacking the essential features)
Prediction	Anticipates unknown data	Conditions of application unclear
Explanation	Uses plausible mechanisms to match outcome data in a well-defined manner	Model is brittle, so minor changes in the set-up result in bad fit to explained data; bugs in the code
Theoretical exposition	Systematically maps out or establishes the consequences of some mechanisms	Bugs in the code; inadequate coverage of possibilities
Illustration	Shows an idea clearly as a particular example	Over interpretation to make theoretical or empirical claims; vagueness
Analogy	Provides a way of thinking about something; gives insights	Taking it seriously for any other purpose

The End

Bruce Edmonds: bruce@edmonds.name

Centre for Policy Modelling: http://cfpm.org

A full paper on this is at:

http://jasss.soc.surrey.ac.uk/22/3/6.html

These slides are at: http://cfpm.org;/slides

More about pitfalls: http://cfpm.org/discussionpapers/236

Some Pitfalls in When Using ABM for Policy Issues

Promising too much

- Modellers are in a position to see the potential of their work, and so can tantalise others by suggesting possible/future uses (e.g. in the conclusions of papers or grant applications)
- They are tempted to suggest they can 'predict', 'evaluate the impact of alternative polices' etc.
- Especially with complex situations (that ABM is useful for) this is simply deceptive
- '*Giving a prediction to a policy maker is like giving a sharp knife to a child*'

The inherent plausibility of ABMs

- Due to the way ABMs map onto reality in a common-sense manner (e.g. people⇔agents)…
- …visualisations of what is happening can be readily interpretted by non-modellers
- and hence given much greater credence than they warrant (i.e. the extent of their validation)
- It is thus relatively easy to persuade using a good ABM and visualisation
- Only we know how fragile they are, and need to be especially careful about suggesting otherwise

Model Spread

- On of the big advantages of formal models is that they can be passed around to be checked, played with, extended, used etc.
- However once a model is out there, it might get used for different purposes than intended
- e.g. the Black-Scholes model of derivative pricing
- Try to ensure a released model is packaged with documentation that warns of its uses and limitations

Narrowing the evidential base

- The case of the Newfoundland cod, indicates how models can work to constrain the evidence base, therefore limiting decision making
- If a model is considered authoritative, then the data it uses and produces can sideline other sources of evidence
- Using a model rather than measuring lots of stuff is cheap, but with obvious dangers
- Try to ensure models are used to widen the possibilities considered, rather than limit them

Other/General Pitfalls

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

- Context matters!
- In each context there will be many conditions/ assumptions we are not even aware of
- A model designed in one context may fail for subtle reasons in another (e.g. different ontology)
- Models generally need re-testing, re-validating and often re-developing in new contexts

What models cannot reasonably do

- Many questions are beyond the realm of models and modellers but are essentially
  - ethical
  - political
  - social
  - semantic
  - symbolic
- Applying models to these (outside the walls of our academic asylum) can confuse and distract

A false sense of security

- If the outcomes of a model give a false sense of certainly about outcomes then a model can be *worse than useless; positively damaging to policy*
- Better to err on the side of caution and say there is not good model in this case
- Even if you are optimistic for a particular model
- Distinction here between probabilistic and possibilistic views

Not more facts, but values!

- Sometimes it is not facts and projections that are the issue but values
- However good models are, the ‘engineering’ approach to policy (enumerate policies, predict impact of each, choose best policy) might be inappropriate
- Modellers caught on the wrong side of history may be blamed even though they were just doing the technical parts

The uncertainty is too great

- Required reliability of outcome values is too low for purpose
- Can be due to data or model reasons
- Radical uncertainty is when its not a question of degree but the situation might fundamentally change or be different from the model
- Error estimation is only valid in absence of radical uncertainly (which is not the case in almost all ecological, technical or social simulations)
- Just got to be honest about this and not only present ‘best case’ results