

UNIVERSITY OF CALGARY

Artificial Policy: Examining the Use of Agent-Based Modelling in Policy Contexts

by

Stefan Mendritzki

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE

INTERDISCIPLINARY GRADUATE PROGRAM

CALGARY, ALBERTA

JUNE, 2010

© Stefan Mendritzki 2010



UNIVERSITY OF
CALGARY

The author of this thesis has granted the University of Calgary a non-exclusive license to reproduce and distribute copies of this thesis to users of the University of Calgary Archives.

Copyright remains with the author.

Theses and dissertations available in the University of Calgary Institutional Repository are solely for the purpose of private study and research. They may not be copied or reproduced, except as permitted by copyright laws, without written authority of the copyright owner. Any commercial use or re-publication is strictly prohibited.

The original Partial Copyright License attesting to these terms and signed by the author of this thesis may be found in the original print version of the thesis, held by the University of Calgary Archives.

Please contact the University of Calgary Archives for further information:

E-mail: uarc@ucalgary.ca

Telephone: (403) 220-7271

Website: <http://archives.ucalgary.ca>

Abstract

Recent economic events have amplified the importance of considering the appropriateness of different modelling traditions as the basis of policy advice. In the context of innovation policy, the appropriateness of agent-based modelling is explored relative to alternative modelling approaches. Validation, which is the process of evaluating the empirical adequacy of models, is raised as an important but complex factor in evaluating policy models. Analysis of validation generally supports claims that descriptive models are subject to the strongest validity tests. However, the validity-centric approach is criticized as creating models of insufficient scope to be useful for policy. Mid level policy modelling is suggested as an underexploited niche which allows sufficient scope to be useful in policy contexts while still maintaining links to empirical methods. This analysis confronts the idea that an increased focus on the validation of models is all that is required to better support policy decisions.

Acknowledgements

Completing this thesis has been quite a journey and there are many people deserving acknowledgement. First and foremost, thanks to my friends, my family, and my partner, Lindsay Penner. Your encouragement and support has kept me on track. Thanks to my committee, and especially Drs. Richard Hawkins and Cooper Langford, for their patience with my many changes of direction. Allowing me the opportunity to participate in their research projects was much appreciated. Their experience in all things academic will surely be an invaluable resource. Thanks to my colleagues for their support. Special thanks to Dan Meeking and Ray op'tLand for their insights and mutual support. Thanks also to all of the members of the InnoLab for their interesting insights, both towards this project and to my understanding of the topic of innovation in general. Thanks to the Interdisciplinary Graduate Program for offering a truly unique graduate experience. Special thanks to Pauline Fisk for putting up with my many questions. Such an enthusiastic and competent administrator is a boon to the sometimes confusing world of an IGP student. Thanks also to all of my many friends and colleagues who I have had the chance to serve with on the Graduate Students' Association. It was truly inspiring to be surrounded by a group of such committed and passionate individuals. Finally, I would like to acknowledge the funders of my project: the Interdisciplinary Graduate Program, the Government of Alberta, and the Social Sciences and Humanities Research Council of Canada (SSHRC).

Table of Contents

Approval Page.....	ii
Abstract.....	iii
Acknowledgements.....	iv
Table of Contents.....	v
 CHAPTER 1: INTRODUCTION.....	 1
 CHAPTER 2: AGENT-BASED MODELLING AND POLICY.....	 7
Agent-Based Models in Evolutionary Economics.....	8
Trends in Agent-Based Modelling.....	15
The Policy Debate: Optimality and Evidence.....	19
Agent-based Models and Policy.....	25
Two Approaches to Agent-Based Policy Modelling.....	27
Agent-Based Policy Modelling Case Study: Electricity Markets.....	32
Descriptive Modelling and Policy.....	43
 CHAPTER 3: VALIDATION OF AGENT-BASED MODELS.....	 45
Validation and its Role in the Modelling Process.....	47
Multiplying Intermediaries.....	55
Types of validation.....	58
The Limits of Output Validation.....	62
Validity as Relative and Purpose Dependant.....	64
A Working Definition of Validation.....	65
Heterogeneity of Agent-Based Models.....	67
Models as Mediators.....	68
A Taxonomy of Agent-based Models.....	70
Conclusions.....	77
 CHAPTER 4: VALIDITY AND RELEVANCE TO POLICY.....	 80
The Local Tool Critique of Descriptivity.....	81
Implications of Directness of Representation.....	89
The Case for Mid Level Policy Models.....	94
Impact on Defining Mid Level.....	95
The Case Against Directly Specified Descriptive Models.....	97
Rethinking Relevance to Policy Contexts.....	98
Case Study: Mid Level Modelling.....	100
Implications of the Mid Level Shift.....	106
A Way Forward for Incrementalist Policy?.....	107
Beyond Familiarity.....	108
Validity effects.....	109
Normalization of Expectations of Policy Makers.....	110
 CHAPTER 5: DISCUSSION AND CONCLUSIONS.....	 112
Artificial Policy: Development and Prospects.....	114

Role of Descriptive and Abstract Models	115
Status of output of mid level policy models	117
Contribution to the evidence-based policy debate	118
Future Work	119
Prognosis.....	120
BIBLIOGRAPHY	122

CHAPTER 1: INTRODUCTION

The academic policy analysis community aims to provide advice to decision makers on prudent courses of action. In order to support their claims that a particular course of action is desirable, possible, and/or necessary, formal models are often used. This is seen prominently in the area of economic policy making, often with the goal of exploring the systemic effects of policy modifications. Recent economic events have highlighted the extent to which policy advice based upon these models may lead to unpredicted and adverse events. In light of these failures, questions about the appropriateness of different modelling methods have been amplified. Given the substantial reliance upon models to generate economic policy advice, there is significant motivation for policy communities to engage with the ongoing debates about modelling methodology.

To explore these methodological issues, inquiry is restricted on two fronts. First, we are interested only in the impacts of modelling choice on provision of economic policy advice. The primary interest in models is in how they can improve decision making. Second, rather than exploring the rather vast subject of economic policy, inquiry is focussed upon the topic of innovation policy. This is an important sub-domain of economic policy which relates to the frequency, quality, and socioeconomic impacts of innovative activity. In terms of innovation policy, two main supporting streams exist. The status quo is the neoclassical economic approach to supporting policy using models. These models are abstract representations of economic processes in the tradition of equilibrium analysis.

These models have been criticized as being too distant from actual processes of innovation to be useful for supporting intervention in real-world systems. An alternative stream has emerged in the evolutionary economic approach. This approach, derived by analogy with biological evolution, frames economic processes as non-stationary. Rather, of interest are open-ended evolutionary processes which are characterized by continual novelty generation. To explore these dynamics, evolutionary approaches utilize computer simulation models to capture more complex dynamics of innovation. Notable among these simulation methods is the agent-based modelling approach. Proponents of this approach in evolutionary economics claim that agent-based models make more realistic assumptions than their neoclassical counterparts.

An important claim of agent-based modelling is that increased reference to empirical contexts renders such models as more appropriate as a basis of policy. However, this claim has not been systemically analyzed; it is rather proposed as self-evident. The remainder of this thesis will be dedicated to exploring the relationship between the empirical content of models and their appropriateness to policy contexts. The driving research question is: to what extent does agent-based modelling represent a more appropriate tool for generating policy advice than other modelling approaches? On the basis of what criterion are models considered appropriate to policy? The approach to answering these questions is novel in providing an explicit account of both the policy process and a specific modelling practice.

Engagement with methodological debates around appropriateness of different modelling practices is not merely an abstract prospect. Partisans of some modelling methodologies, notably complexity scientists in North America, have attempted to use the current economic issues as a motivator to bring to the fore debates about the effectiveness of prevailing economic modelling practices in supporting policy. An example of this movement is Brown et al (2008), who propose that the economic downturn can largely be blamed on the unrealistic models which were supporting policy decisions. Essentially, it is argued that the models utilized by neoclassical economists contributed to the speculative trading of derivatives.¹ In order to remedy this, Weinstein has called for an “Economic Manhattan Project”, which is, according to Brown et al, a call for a project to establish “a scientific conceptualization of economic theory and modeling that is increasingly reliable.” (ibid, sec. 4)² Many of the arguments utilized in support of the claim that new modelling methods are more reliable are well worn critiques of the prevailing system of neoclassical theorizing and modelling. These are centred on the empirical adequacy of the first-principles assumptions necessary to create neoclassical models. These strong assumptions include market equilibrium, clearing markets, static preferences, extrapolation from the results of 2-party games, perfect information, and a neglect of relative scales. As opposed to these empirical oversights, Brown et al. offer the possibility that their approach will allow us to ~~base~~ the design and regulation of those markets on *correct and verifiable*

¹ The specific explanation posed – that a belief in the objective as opposed to socially consensual nature of the prices of derivatives – may be seen as quite well founded. See MacKenzie (2008) for detailed observation of the construction of market facts. However, our interest here is not in the specific explanation but rather in the evaluation of the role models might play in building an alternative approach.

² The implicit implication is the substantial resources which would be needed to support such a project.

principles and models.” (emphasis added, *ibid*, sec. 3) They claim that their modelling approach can offer strong support to policy makers:

The role of an independent, non-partisan scientific conceptualization of economics should be to provide these policy makers with a notion as to the likelihood that new economic and financial regulations they are considering will have the results they desire and that these will not involve unintended consequences to others. (*ibid*, sec. 4)

Here we see the ambition of the claim: a change in modelling methodology is the key factor in the provision of more reliable policy advice. The methodological shift is premised on creation of more empirically grounded models. More broadly, these arguments attempt to utilize empirical evidence to remove “~~opinion~~ and ideology” from policy, rendering the conduct of policy advice a domain for objective scientific inquiry. The claims of Brown et al. link to the claims of the broad movement of evidence-based policy, in that both argue that increased reliability of decisions making is based on increased empirical grounding. To explore these issues, we turn to outline the argument which will be made. Each chapter is motivated by a focussed question which is an important part of the argument at hand.

In Chapter 2, we will ask: what is the current state of the art in terms of agent-based models broadly and their application to policy questions specifically? A preliminary step is to

provide a fuller explanation of the relationship between evolutionary economics and agent-based models. This will lead to a discussion of the nature of agent-based modelling, with a focus on the advantages proponents claim for it and current trends. Two of these trends will be of primary interest: the turn towards the policy applications and an open debate on how to empirically validate such models. These issues will be made more concrete through a case study of policy modelling in the case of the exercise of market power in electricity markets. This will lead to a discussion of recent attempts at providing guidelines for policy modelling which claim that the key to more appropriate models is better validation.

Chapter 3 turns to examination of empirical validation of agent-based models. How is the empirical validity of models determined? Can we say anything general about the relationship between modelling methodology and their validity claims? It is argued that this area lacks definitive literature which can be utilized for our purposes. Therefore, an attempt is made to provide an overview of the major streams of the relevant literature. This aims to both highlight open issues in the validation process and to enable the formation of a well founded working definition. Having developed such a definition, the implicit assumption of the validation literature of a homogeneous class of agent-based models will be challenged. In highlighting the significant heterogeneity in agent-based models, a key dimension will be argued to be the balance of theoretical and empirical elements of which a model is composed. This difference in model type will be explored in terms of its implications for the prospects of evaluating validity. Models which are formed on the basis of more empirical elements will be argued to be subject to much stronger validation tests

(and thus claims) than models of more theoretical composition. This analysis calls into question the strength of some of the validity tests proposed by evolutionary economists.

Chapter 4 returns, validation framework in hand, to the problem of using agent-based models to generate policy recommendations. Is consideration of validity the primary judgement to be made about policy models? How does this relate to the evidence-based policy approach? Based upon arguments about validity, it is argued that if validity is the primary evaluative criterion then the evidence-based policy approach is well founded. However, in examining the practice of descriptive agent-based modelling, the lack of descriptive models outside of highly localized domains will be argued to be a major issue. Through the use of example models, a general argument about the prospects of developing descriptive models of non-local domains will be explored. The implication of this will be seen to be an argument in support of mid level modelling for policy. Benefits of adopting a mid level approach are explored both from application to a case study and as it relates to ideas of evolutionary or incrementalist policy making.

Chapter 5 concludes by examining the implications of the arguments presented for the research questions. An important area of research is suggested to be a reorientation of expectations of the impact of models generally on policy contexts. Wider implications in terms of the evidence-based policy movement are drawn. Finally, directions for future research are discussed.

CHAPTER 2: AGENT-BASED MODELLING AND POLICY

The core of this thesis is the relationship between the use of agent-based modelling methods in evolutionary economics and the possible contribution of these methods to policy making.³ Before analysing this relationship, it is relevant to provide some background into the domains which are being related. The introduction of agent-based models in evolutionary economics will be framed relative to the original motivations for turning to simulation methods, mainly to escape the constraints of previous equilibrium models. It will then be claimed that the current discourse around agent-based modelling in evolutionary economics reflects development of similar themes. Having examined the claimed advantages and challenges of agent-based modelling, emerging literature on trends in methodology is reviewed. We then turn to policy, giving some background on the different theories of policy analysis. This will focus on the optimal policy tradition associated with neoclassical approaches and three traditions which have grown in contrast to this position: evidence-based policy, incrementalist policy, and evolutionary policy. With this background in hand, we will proceed to examine the literature on agent-based policy modelling. To demonstrate the difference that turning to agent-based modelling makes to policy recommendations, a case study modelling electricity markets will be explored. This leads into discussions about the direction of the policy agent-based modelling debate and the importance of validation.

³ While the specific context under study is evolutionary economics, in considering methodological questions, the scope shall be wide enough to take into account methods in different branches of social science.

Agent-Based Models in Evolutionary Economics

The use of agent-based modelling in evolutionary economics is a relatively recent development. Agent-based modelling is a computer simulation methodology which explores the emergent relationships between software agents through decentralized evolutionary processes. It allows for the explicit representation of the dynamics of evolution and highly flexible application to a wide range of situations. Before turning to recent literature on the topic, it is relevant to first examine the origins of the use of simulation methods in evolutionary economics.

Nelson and Winter (1982) provided the foundational piece on motivating factors behind the use of computer simulation methods in evolutionary economics. They argued that within the orthodox tradition of economics, the process of research had yielded successively more sophisticated mathematical tools but little accompanying conceptual or theoretical development. In turning from the models of neoclassical theory (hereafter analytical models, following the terminology of Nelson and Winter) which were based upon the concept of equilibrium, it was hoped that greater progress could be made without sacrificing rigour. By freeing research programmes from the constraints of the strong assumptions of the analytic framework, researchers were enabled to create models based upon different assumptions. The classical case of this is the argument for modifying agent decision functions from those of Rational Action Theory (RAT) to those of Simon's (1996)

bounded rationality.⁴ Calculus required assumptions of homogeneity of agent characteristics and behaviours in order to produce tractable models. The problem with breaking from these assumptions was that while RAT had an obvious operational definition in formal models, bounded rationality did not. Computer simulations offered a modelling approach which was still formally defined, yet less constrictive.

Though Nelson and Winter argued for simulation methods, they did not advocate for agent-based models specifically. In fact, agent-based modelling was not well developed as a method at the time of their simulationist turn. The simulation models utilized by Nelson and Winter attempted to stay close to the traditional modelling goal of tractability. As links expanded to evolutionary modelling more generally, it was found that there are several different types of simulation methods which might be utilized to support evolutionary explanation. For example, one might turn to replicator dynamics, which is a population based approach. However, replicator dynamics models remain quite abstract. It is difficult, upon detailed consideration, to link elements of such explanations to specific empirical contexts. This is another important point of the evolutionary school: a rhetoric of making more realistic assumptions about economic mechanisms than neoclassical theory. Agent-based models present opportunities on this front, in that they facilitate reasoning

⁴ In neoclassical models, economic agents were considered to be reducible to simple maximizing functions. Given information inputs (generally global, another critique), the agents would compute the maximum output. This approach was criticized by Simon who proposed that the assumptions of the economic models were untenable on psychological grounds. That is, no human cognition process could possibly do the calculation assumed by the economic models. Further, looking at humans in actual decision making contexts, their behaviour could be observed to differ substantially from the RAT assumption. Instead of maximizing outcomes, agents were seen to use heuristics (rules of thumb) to generate good but not optimal solutions. Since this process aimed for satisfactory and not optimal solutions, it was termed satisficing. The general framework became known as bounded rationality.

about the dynamics of evolution as opposed to only the outcomes. They offer a way to define agents at one level whose dynamic interactions result in emergent results at a higher level of aggregation. For example, by defining a model where agents are firms, the emergent result of their interactions can be interpreted as the dynamics of industries. Analytic models explain by similarity of modelled and observed data sets, by the congruence of data models. There is no claim that the data sets are produced by similar mechanisms, meaning that models are not evaluated on the empirical accuracy of their generating processes. Model assumptions which are clearly at odds with empirical results are thus acceptable (for example, the assumptions of RAT). As opposed to analytic models which only investigate outcomes, agent-based models allow exploration of dynamics.⁵ Since it demonstrates directly the dynamics which generate emergent patterns, it has been termed generative explanation (Epstein, 2006). Given the interest of such models in process-based accounts, their form of explanation could also be termed algorithmic explanation. A focus on underlying processes allows exploration dynamic, non-equilibrium patterns unfolding through time. This makes agent-based modelling relevant to evolutionary economics. The claim being made is not that all evolutionary economic models are agent-based, but that a motivation exists to take them up.

Themes of flexibility of assumptions and the related ability to become closer to empirical contexts are seen in recent proposals for the possibilities of agent-based modelling. Pyka

⁵ The idea of exploring only outputs vs. exploring how outputs are generated has been captured in ideas of black-box vs. white-box models. Black-box models are those which have no interest in the mechanisms which produce outputs but only that certain inputs lead to certain outputs. If the mechanisms of interest are open to investigation, then the model is a white-box (sometimes also called glass-box).

and Fagiolo (2005) examine agent based modelling in the context of neo-Schumpeterian economics. They claim it can be primarily seen as a step towards descriptive adequacy, in that it allows modelling of more realistic assumptions than analytical models. Pyka and Fagiolo summarize the key assumptions of agent-based models as autonomous agents, heterogeneity of agent characteristics, bounded-rationality, true uncertainty, locally networked interactions, and non-reversible dynamics. These are formed in contrast to analytical models which feature identical and perfectly rational agents (the ‘representative agent’), which are not represented individually, interact only through abstract market signals, and come to the same equilibrium regardless of initial conditions. We take up each of the assumptions of agent-based models in turn. The autonomous agents assumption is that agents are individually represented and make decisions based upon individual characteristics and histories. The individual identity of agents allows their characteristics to differ both within and between models. This heterogeneity of characteristics may derive from differing initial conditions or different histories built up within the model. The bounded-rationality assumption (discussed above) along with the true uncertainty assumption result from relaxing the requirement of perfect rationality. The heuristic (rule-of-thumb) decision making procedures of bounded rationality inherently exclude a global understanding of opportunities and risks. This, in turn, means that agents are not in the risk taking framework of analytical models, where risks are calculated and known, but the true uncertainty of not even knowing what the risks are. The non-trivial interaction structure assumption refers to the fact that agents interact and communicate directly rather than just through abstract market signals. These networks of interaction and communication are

often quite complex. Finally, the non-reversible dynamics assumption refers to the historical contingency of such systems due to the variety of agents and their complex interactions. Since actions at time t depend on the state of the system at time $t-1$, actions at any point in time are predicated on historical causal chains. This is captured by the terms non-reversible dynamics (time has a direction) and path dependence (current agent behaviours depend on the historical path to the present). Such evolutionary processes are open-ended, as there is no final state which signals the end of the simulations. There is no guarantee that a system will settle to equilibrium or, if it does, that it will settle to the same equilibrium from different initial conditions.

The relaxed assumptions of agent-based models offer the flexibility to capture a greater complexity of dynamics and improve the plausibility of the assumptions. The innovation of agent-based modelling is in allowing localized and parallel interaction, enabling richer representations of dynamics (for example, path dependence). This richness of representation allows exploration of collective phenomena (emergent behaviour) under a variety of different sets of assumptions. Flexibility in model assumptions and structure are also intended to provide greater scope for validation against “real-world observations”. This last point is linked to the greater congruence of assumptions with observations than seen in analytical models. The other side of the realism coin is that increasingly “realist” assumptions⁶ translate into more difficulty in studying the system and decreasing clarity in

⁶ Such as increasing agent heterogeneity (e.g. increasingly sophisticated boundedly rational behaviours), allowance of structural change within the model (structural innovations), etc.

the causal connections between assumptions and outcomes. The very stylization of analytical models made them amenable to providing precise conclusions.

This provides a fair overview of what is meant by agent-based modelling and the trade-off which it presents. This explanation suffers to an extent as an explanation of agent-based modelling of social phenomena more generally, in that many of the claims being made are in comparison to the neo-classical economic tradition. It is probably a fair evaluation to say that Pyka and Fagiolo effectively isolated the key features of agent-based modelling in social contexts more generally but their judgements on how these features are implemented is quite couched in the disciplinary standards of economics. Take bounded rationality for example. In almost all agent-based models there is a break from the tradition of assuming the sort of omnipotent and omniscient qualities seen in rational choice models. The open issue is what the implications of this negative heuristic are. Is agent cognition still essentially “rational” (a difference of degree) or is it something entirely different (i.e. imitative)?⁷ This applies to many of the other descriptors of agent-based modelling which are negative correlates of constraints considered in traditional modelling. For example, the idea of connecting agents through a local network raises issues of which type of network. While the assumptions of agent-based models move away from the strong definitions of analytical models, it is an open question as to how far and on what basis they will diverge.

⁷ The algorithmic flexibility of agent-based models raises the possibility of utilizing results from the behavioural schools of economics (such as neuro-economics, experimental economics, and behavioural economics) in the model development process.

Another point which deserves attention is the various traditions which could be classed under what has been broadly termed agent-based modelling. Pyka and Fagiolo note that the subject matter they are addressing themselves to might self identify as evolutionary economics, agent-based computational economics (ACE), neo-Schumpeterian models or history-friendly models. What they fail to mention is that the same can be applied to the methodology. Variants include complex adaptive systems (CAS), multi-agent systems (MAS), social simulation, agent based social simulation (ABSS), evolutionary game theory, artificial societies, companion modelling, genetic algorithm (GA), cellular automata, etc. In both cases, there are threads of commonality but also real differences of approach.⁸ In the case of the difference economic schools, one can expect a more formal, more generalizability-focused treatment in a neo-Schumpeterian model than in a history friendly one. Methodological designates follow a similar pattern. Companion modelling is tied to a specific, small scale case; artificial society models are quite foundationalist; game theoretic models incorporate highly structured decision situations; GA models incorporate strong biological metaphors about information storage, variation, and selection; MAS derives from distributed artificial intelligence and is generally quite abstract. As with any terminological issue, where the waters are cloudy it is imperative not so much to declare allegiance as to provide explicit definitions. Here the term “agent-based modelling” is adopted as a general term to specify modelling traditions, incorporating (or at least

⁸ In some sense pretending that the difference does not exist exacerbates the problem. One may read in the abstract of a paper that it is about “agent-based modelling” only to find half way through that the author meant this as synonymous with a genetic algorithm model or an iterated, evolutionary prisoner’s dilemma model.

compatible with⁹) simulation construction, where agents with autonomy and heterogeneous characteristics jointly determine emergent outcomes through locally networked interactions. This is aligned with Pyka and Fagiolo's terminology. Subclasses of this high-level designation may be identified by additional constraints which they impose. We will also adopt the standard terminology of calling the referent of the modelling process the target system. Further classification will be addressed by adoption of an agent-based model taxonomy in Chapter 3.

Trends in Agent-Based Modelling

Being a relatively new area of study and being fragmented in diverse disciplinary contexts, agent-based modelling has largely lacked common definitions and standards. However, there have been some very recent developments in the methodological literature. Some recent reviews are taken up to explore the state and future direction of agent-based modelling.

Heath et al. (2009) suggest that agent-based modelling should be seen as a method in development; it shows great potential to deliver powerful results but has delivered few concrete results to date. In their literature survey (of 279 articles published between 1998

⁹ A simulation model is a formalized model, which, as explained in Chapter 3, is the outcome of a set of intermediary models. Thus we might consider something like Windrum and Garcia-Goni (2008), which specifies a model but does not endeavour to simulate that model, to be within the tradition despite a lack of simulation. In the usage here, modelling and simulation will generally be taken as synonymous, unless it bears explaining how they differ.

and 2008), an important emphasis is on the criteria and techniques for validation of agent-based models. They separate validation into two types: operational (comparing model outputs to actual outputs) and conceptual (comparing the conceptual model to the somewhat ambiguous “system theories and assumptions”). According to these definitions, approximately 1/3 of models had no validation, 1/3 had either conceptual or output validation and 1/3 both. Interestingly these results were highly time dependant. For example, at the beginning of the period 75% of articles featured no validation versus just 12% at the end of the period. Clearly, validation of some form is becoming standard practice.

The authors developed a taxonomy of the purpose of simulation models, based upon level of understanding of actual systems that the models are meant to provide. These were: generator (models whose primary purpose was hypothesis generation), mediator (models whose purpose is to provide insight but not full understanding of a system), and predictor (models whose primary purpose was predicting specific outcomes). According to the definitions of the survey, there were no models in the sample which fell into the predictor category. Models in theoretical categories (social sciences and economics) tended to be more focused on generator models, while applied domains (business and policy) leaned towards a mediator approach. It could be argued that the taxonomy employed is flawed due to the Predictor category assuming a level of predictivity which is rarely claimed in practice. However, it is useful that such a strong definition was used in that it puts in its place the literature which takes prediction (in the strong sense of point prediction) to be the

goal of modelling. There is little evidence that this is the direction the community has been going. A final point made is that propensity to validate is related to the purpose of modelling. Models in the mediator category (i.e. closer to understanding real systems), were seen to employ validation more often.

Another relevant review is Meyer et al (2009) who report the results of a bibliometric study on the first 10 years of an important journal in the area, the Journal of Artificial Societies and Social Simulation (JASSS). They compare the co-citation patterns in the periods 1998-2002 and 2002-2007. Based upon this, they conclude that the typical signs of an increasingly mature disciplinary structure are observed (i.e. movement from generalized concerns and justification of the method to specific topical clusters; movement from foundational books to dedicated journals). This result suggests that agent-based modelling has a sufficiently developed tradition and community to give plausibility to its contention as a legitimate source of knowledge. As will be seen, this notion of legitimacy (both internally within the academic community and externally as a basis of action in the world) is prevalent in research on agent-based modelling method.

Finally, Pyka & Werker (2009), in the introduction to a special section on methodology in agent-based modelling, identify the additional flexibility allowed by agent-based modelling as a methodological challenge. Due to a lack of constraints typical of the analytical approach, agent-based modelling underwent an explosion of methodological plurality (or, less generously, methodological chaos). This observation is not novel, as other calls exist

for strengthening of methodological standards in agent-based modelling (e.g. Richiardi et al, 2006). The issue is that the previous analytical models were undertaken within the context of what might be called *methodological social modelling*¹⁰, where methodological standards were well defined and explicit. This, in turn, greatly facilitated judgements of legitimacy and encouraged cumulativeness of models. A good model was most commonly one which took an existing model and modified it slightly. With these standards in hand, a model could be judged internally (within a discipline) and externally (as commendable to outsiders as reliable). Given that one of the goals of the turn to simulation was to escape such methodological conventionalism, it is not surprising that strong guidelines have been difficult to come by, especially given the diversity of modelling goals and implicit assumptions. Implicitly recognizing this, Pyka and Werker avoid proposing comprehensive guidelines, and instead suggest three methodological priorities. First, the discipline must do a better job of communicating with computer scientists on acceptable practice and emerging findings in computation. Second, more discussion around the topics of model setup and validation is required. Third, investigation into the use of simulations for consulting with policy makers is required. We might think of the second and third as internal and external legitimacy respectively.

The shift towards the epistemological, methodological, and policy aspects of models is generally evident in recent literature. For example, special issues of Computational Economics (Fagiolo et al., 2007) and the Journal of Economic Behaviour and Organization

¹⁰ Methodological social sciences draws on the concept of the methodological social sciences (psychology, economics, etc.), where highly detailed standards of methodology have been developed.

(Dawid & Fagiolo, 2008) have addressed the issues of validation and policy modelling respectively. Given increasing maturity of method, the thesis that agent-based modelling is poised to have an impact on policy analysis is sufficiently plausible to motivate further research.

The Policy Debate: Optimality and Evidence

In order to understand the effect on policy of a change in modelling methodology, we first need some idea of the different streams of thought around possible academic input to policy. Here we will start by presenting critiques of what is taken as the status quo of academic input to economic policy: the optimal policy stance. Having identified issues with this approach, it will be suggested that there are three prominent ways out of the idea of optimal policy. One option is evidence-based policy, which aims to be the antithesis of the theory-based optimal policy stance. The other two, evolutionary policy and incrementalist policy, are less clearly placed within the empirics-theory discourse. These framing approaches will allow us to suggest in the next section the possible directions for agent-based policy modelling.

If we return to Nelson and Winter (1982), it is notable that part of the push towards the evolutionary approach was a motive to be more useful to policy contexts. Their critique of ‘orthodox’ economics and its role in policy making was centered on analyzing the prospects of the concept of optimal policy. Economists traditionally took the development

of policy, as any other problem, to be one of optimization. Roughly speaking, the process began with the model defined within the disciplinary context (the internal/scientific model). This defined the theoretical optimum at which an economic system could operate. Empirical results from actual economies would then be compared to these theoretically derived results. Divergence from the optimum led to recommendations that making the conditions in the target system closer to those in the theoretical model would move outcomes closer to the theoretical optimum. So the optimal policy approach establishes an optimal state and suggests mechanisms to move towards this state. Nelson and Winter argued that this runs into several issues in the evolutionary framework. First, it is not clear that it is ever possible to identify the optimum in multi-factor economic systems, let alone achieve it. Second, this view ignores the historical development of firms involved and thus is likely to fail. Third, the options considered are theoretical and not practicable; trying to maximize a theoretical indicator is unlikely to produce desired practical effects. Instead, they suggest that a more fruitful approach is to move to a more relative frame of evaluation. First, researchers should attempt to establish issues with the current system. These may be developed in reference to the operation of other comparable systems (i.e. a different jurisdiction) but are always relative. Second a set of well defined and practicable alternatives should be developed. Finally, the alternatives should be compared to the current operation of the system to determine the best alternative in the given context. What is notable at this point is that policy solutions are developed at the level of the problems which motivate them. This is a key point which we will return to. These ideas on the

inadequacy of the optimal policy approach are similar to the earlier work of Lindblom, which we turn to now.

Lindblom (1959, 1979) argues that we can think of two modes of policy studies. The first, he calls synoptic, which signifies an attempt at comprehensive analysis from an objective viewpoint. Such a view proceeds by first identifying desired ends and then analysing means which best achieve these. To the extent that the objectives are properly specified and the means properly analyzed, the resulting advice is considered the best course of action. This is just the optimal policy paradigm of Nelson and Winter. The key debates are then whether the ends are correctly specified and, more often, whether specific courses of action will in fact deliver the ends which they suppose. This tends to be an abstract debate in terms of results of general approaches in highly structured problem domains. Lindblom suggested that this form of analysis was dominant in policy circles.

An alternative, which was argued to be the norm in actual policy making, is the incrementalist approach. The foundation of this approach is management of complexity through moving only incrementally from current patterns of action. By starting the process at the status quo, knowledge of complex interactions can be leveraged. In terms of process, incrementalist analysis begins by clarifying perceptions of the status quo. Next, incremental alternative scenarios are developed. These are then subjected to comparative analysis to understand the differences in outcomes which they might bring, the so-called marginal comparison step. Then just the marginal differences are weighed in terms of

some value system (i.e. how much prosperity is worth how much equality). A choice is made and implemented, becoming the new starting point for analysis. This is again quite similar to Nelson and Winter. It is interesting to note parallels in terms of satisficing and optimizing in areas such as computer program design (linear vs. agile methods), evolutionary search models (local vs. global search), and theories about entrepreneurship (effectuation vs. analysis).

Lindblom (1979) notes that the incrementalist approach has been subject to several criticisms as a method of developing policy. First, that the incrementalist approach is inherently conservative. Second, that it is subject to issues of momentum and path dependence, possibly leading to adverse effects. Third, that it has the effect of separating issues which are systemically linked.¹¹ Lindblom responds to these critiques along two paths. In the first place, the incrementalist theory can be seen from the perspective of being descriptive as opposed to normative.¹² In describing the actual process of policy making, the issues with the prevailing synoptic or optimal policy paradigm are highlighted. In addition, it is exactly by recognizing how synoptic policy advice fails to have impact that improved forms of analysis may be discovered. This is the source of the strategic continuum in Lindblom (1979). Here, incremental and synoptic arguments are the poles of a continuum. The normative argument put forward is that moving beyond the purely incremental may lead to better informed decision making but that consideration of impact

¹¹ We might think of part of the problem with homelessness policy is treating it as independently of housing, education, and minimum wage policy.

¹² A provocative interpretation is that Lindblom did for policy analysis what Veblen did for economic analysis.

should temper the desire to go too far on the continuum. In trying to do better than incrementalist analysis, we might maintain a useful baseline and avoid falling into an optimal policy mindset.

In order to move beyond the optimal policy paradigm, incrementalism is only one option. A view that shares this focus on steering from the status quo is the evolutionary approach offered by Arthur (1999). A prominent complexity theorist, best known for modelling of increasing returns, Arthur discusses the shift from a static to a process oriented view. In traditional forms of economic modelling, the structure is specified with an eye to achieving some outcome. Instead, he argues that economies are places of perpetual evolution and novelty. In such a system, it no longer makes sense to think of static outcomes which are being worked towards. Rather than specifying the ends, policy should intervene to shape the emerging outcomes in a desirable manner. Specifically, it should “seek to push the system gently toward favoured structures that can grow and emerge naturally.” (ibid, p. 108) This leaves many questions unanswered. How exactly does one “push the system gently”? What is a “favoured structure”? Favoured by whom and on what basis? How do we know that the relevant structures are in fact economic in the first place? Evolutionary theory at this level is an interesting prospect but not a practicable approach. Perhaps policies do “succeed better by influencing the natural processes of formation of economic structures than by forcing static outcomes.” (ibid, p. 109) but practicing this approach requires evaluation of the extent to which we, currently and in theory, can know the natural processes in order to influence them. It is also an approach which does not give explicit

status to the decision making context. While evolutionary at the level of the target system, it is not evolutionary at the level of policy making.

A final school of thought is the evidence-based policy approach. It has become prominent in areas such as health policy and corrections policy (associated groups are the Cochrane Collaboration and the Campbell Collaboration respectively, see Petrosino et al., 2001 for an overview). The basic argument put forward is that more systematic studies of efficacy of different approaches are the key to better political decisions. This model proposes that scientists collect evidence which is then reported to decision makers. This assumes a direct influence of research evidence on the policy process. An implicit assumption of this program is that the evidence adduced represents reliable knowledge, as opposed to the opinion, conjecture, and ideology of theory. One criticism of this approach is the turn of phrase of Borden and Epstein (2006), who suggest that the reality of trying to implement this process is policy-based evidence. This is the tendency to select policies first and then rationalize them in terms of research which is designed to validate the premises. While the pathways for advice are relatively clear in terms of restricted questions (the safety of a particular drug, for example), there is not much theory beyond this of how the evidence base transfers to the policy context.

These different modes of policy analysis set up very different relationships between academic research and policy, which in turn provide clarification of the possible role of policy modelling. In the optimal policy view, the role of policy modelling is to take

different sets of criteria and demonstrate the best way of reaching those criteria. It is then the policy maker's role to set the criteria. The evidence-based view sees the role of policy modelling as the provision of faithful representations of the domain upon which policy is being made. Policy makers then take these representations as starting points in their decision making process. Both the optimal policy and evidence-based policy views see a very direct impact of research on decision making. They both attempt to put their results at the centre of the policy process and have policy makers adjust to their results. In this sense, policy analysis becomes an applied science. The incrementalist policy view provides little scope for application of modelling in its strong form. In the terminology of the strategic continuum, it is possible that modelling may provide an extension of incrementalist analysis. Finally, the evolutionary view suggests that modelling may represent existing structures and possible modes of influencing their unfolding evolution. It is a tempting hypothesis to draw parallels between the strategic extension of the incrementalist approach and the evolutionary approach but this is a complex issue which we do not take up here. This analysis of the differences between different modes of policy advice sets the background for the more specific literature on agent-based modelling and policy.

Agent-based Models and Policy

Using agent-based models to explore policy issues has been considered for some time, though largely as only a theoretical possibility. For example, Pyka and Fagiolo (2005, p.

21), in their definition of agent-based modelling state that “Thanks to the flexibility and the power of agent-based approaches, it is easy to conceive frameworks where policy experiments are carried out to evaluate the effectiveness of different policy measures (e.g., anti-trust policies), for a range of different institutional setups and behavioural rules.”

Despite this possibility, impact has been slow to emerge. In their introduction to a special issue on the use of agent-based models for policy evaluation and design, Dawid & Fagiolo (2008), give the lay of the land. According to their reading of a leading economic modelling handbook on agent-based models (Judd & Tesfatsion, 2006), the majority of work had been devoted to descriptive as opposed to normative work. However, Dawid and Fagiolo point out that detailed policy modelling does exist, such as models of energy markets (Sun & Tesfatsion, 2007), the U.S. coffee market (Midgley et al, 1997) or the pharmaceutical industry (Malerba & Orsenigo, 2002). In engaging the issue, the special issue offers a variety of approaches without much analysis on how they are related. A sample of model topics includes:

- Anti-trust, public support, and public procurement policies on the rate of industry concentration and technological change
- Vertical integration in electricity markets and effects of manager incentives on prices and profits
- Online auction design, a comparison of hard and soft close auctions
- Impact of modification of EU agricultural subsidies for farm structure, prices, and farm profits

Though some of the papers in the issue are still at a rather abstract level (i.e. the ability of Tobin taxes to stabilize theoretical stock exchanges), there is a noticeable turn to specificity. The authors theorize that policy makers might be more willing to heed advice from models which contain familiar structure rather than general insights from rather abstract models. In terms of barriers to adoption, they highlight a dearth of validation and reliability testing standards. To what extent will the results derived in models translate to a reliable basis for policy decisions? The outstanding issue is that this special issue provides little indication about the more general character of agent-based policy modelling. We now turn to recent approaches to this issue.

Two Approaches to Agent-Based Policy Modelling

Up until very recently, the literature lacked explicit accounts of the relationship between policy and agent-based models. The special issue on agent-based modelling methodology mentioned above (Pyka & Werker, 2009), provides two recent attempts at providing a framework for this relationship. The first is an attempt to extend a particular practice within agent-based modelling to the task of advising policy. It is from the perspective of researchers whose background is in modelling. The second differs in that it attempts to lay out requirements of models from the perspective of policy. While the first view sees policy modelling as applied science, the latter views models as tools to support an autonomous domain of study. It will be suggested that the latter approach is much more strongly informed by insights into the policy process.

Brenner and Werker (2009) explore what sort of policy advice might be derived from agent-based models. They are explicit in framing their account of agent-based modelling in a specific philosophical tradition (critical realism) and a specific method (abduction). Very briefly, their abductive modelling approach attempts to use empirical data to restrict model parameters and to determine if a given parameter set replicates observed patterns. The goal of this process is to uncover the structural properties of historical processes. Using this knowledge, an attempt is made to predict a reasonable range of possible outcomes. The details of this framework are not important for our interests here. The point is that the process which is proposed for policy modelling is simply the abductive method which the authors have presented elsewhere (Brenner, 2001) with the addition of a policy implications step. This framework is more interested in promoting a specific modelling framework than in addressing the character of relationships to policy analysis. It is very much in the tradition of seeing policy as applied science rather than an autonomous domain.

This lack of interest in the policy context is seen in a rather naïve view of the role of modelling in informing policy. The important task is seen as the elucidation of the roles of causal economic structure (as opposed to chance). Policy is framed as containing errors which modelling can remedy. The very ambitious goal is to find policies of high effectiveness and efficiency, low risk of failure, and low risk of unintended results. Though a policy study which yielded these results would be very valuable, many policy scholars might question the extent to which any policy ever studied meets these lofty

standards. They might also be quite sceptical as to how we were going to arrive at these results.

The particular example they provide bears out these observations. They ~~address~~ address the question of whether the support of private innovation activities and start-up activities—e.g. by financing a related public research organization—in an industry in a specific region for a limited period of time increases the chances that a local cluster emerges and sustains in this industry and region.” (ibid, para. 4.1) Using a knowledge spillover framework, a high dimensional model is used to include as many empirically relevant parameters as possible. The model is then subjected to the abductive procedure. A baseline is created for the case with no policy support, which yields a probability of a cluster developing. The impact of policy measures is straightforward because it assumes some simple, instrumental effects of policies on parameters and interrogates the results of these instrumental effects. For example, the introduction of a public research institute is represented by multiplying by a factor of 10 both the rate of start-ups and the basic innovation activities of firms. This is then used to investigate the effects of policy on the probability of the emergence of local clusters. Based upon this model, it is concluded that the policy measures discussed increase the likelihood of the emergence of a local cluster most significantly at an early of mature stage (0 or 10 years). Furthermore, it is suggested to focus on industries dependant on external knowledge. These results are both supported by additional empirical literature which points in the same direction.

The question to be asked of this approach is: if they simply replicate the results of empirical studies, what additionality do they offer? It is claimed that ~~this~~ methodology leads us beyond the common use of simulation model, as we are able to infer characteristics of classes of systems that have a general validity and are able to provide valuable advice for policy.” (ibid, para. 5.3) It is not clear what is meant by general validity or how this translates into valuable policy advice. Also, there are questions about the extent to which policy advice can be derived so directly from a rather stylized simulation. The rather instrumental effects of policy changes proposed is a very important part of the model which remains completely conjectural. Such an approach to modelling does not clearly tell policy makers anything about the impacts of public research facilities or even their claimed instrumental effects of increasing start-up rates and basic innovative activities.

An alternative approach which explicitly takes into account the decision context of policy makers is provided by Yucel and van Daalen (2009). They propose a two pronged approach to policy model evaluation. Both the reliability of the model and its relevance to policy making contexts are considered to be important. In support of the latter, they note that the quality of the study does not necessarily imply acceptance of results by clients and/or colleagues.¹³ So we may have a valid model but if it incorporates strategies or

¹³ The process of creating scientifically sound models which are nonetheless rejected by clients is noted in the literature. Alam et al (2007) note an example of this statement in the case of water management studies conducted by Barthelemy (2006). Due to incompatibility of recommendations with existing practice, the developed model was rejected. Similar results were seen in Pahl-Worstl and Hare (2004). However, the rejection of the models is not systemically linked back to a failure to design models in a way which is useful to policy makers.

policies incongruent with the client, the model will not be relevant. The dual part nature of his input is summarized in the processes of building up confidence in a model.

They differentiate potential roles of models according to the Mayer et al (2004) hexagon model of policy analysis roles. Policy analysis may fall into a number of different activities: analysis, design, clarification of arguments, strategic advice, mediation, and participation. The outcomes of these various modes would be: generation of knowledge, evaluate sets of interventions, clarified arguments and values, strategic political advantage, consensus building, and involving interested parties. The primary result of this line of thought is the realization that some activities have goals which require validation (i.e. analysis and design), while others less so (i.e. participation, mediation).

Policy input is accomplished through scenario development for “plausible and relevant conditions” for the target system. For policies considered, it is not enough that they solve the problem at hand; they must also fit existing organizational structure. If structural changes are proposed, evaluation of the feasibility of implementing such changes is required. Finally, when dealing with structural change we must contend with the issues of what Yucel and van Daalen term “sleeping structures”. In any model, boundaries must be drawn as to which structures are important to include and which can be neglected. The insight of sleeping structures is that these boundaries may be dynamic and contingent, shifting when structural change occurs. Thus, a sleeping structure is a structure which is excluded from a model because it is non-central under one set of conditions but which may

be “awakened” by structural change. If the boundary between central and non-central structures is not static, there are significant challenges to analysis of systems which have undergone substantial structural changes. It is no longer clear that regularities which are taken for granted under current conditions will still hold under a new structure. This analysis tends to complement that of Lindblom about the limitations of prediction in the context of structural change. Despite providing an approach which is more subtle from a policy perspective, the analysis of Yucel and van Daalen fails to specify what is unique about *agent-based* policy modelling. That is, their analysis could be read as similar to Lindblom in giving cautionary advice about considering only the scientific merit of models in general.

Agent-Based Policy Modelling Case Study: Electricity Markets

The literature cited above discusses general issues with agent-based policy but fails to provide insight into the difference of turning to agent-based models particularly. In order to shed some light on the situation, it is relevant to examine a case study in which agent-based models can be directly compared to other forms of modelling. This can be framed in terms of the discussion of the relative properties of analytical and algorithmic explanation offered above. The general argument is that the representational flexibility allows representation of specific contexts, which may increase relevance to policy makers.

What makes this difficult is the variety of models which exist under the banner of agent-based policy models. Take, for example, the Dawid and Fagiolo (2008) special issue. Some of the studies are at the level of specific systems and specific changes which might be made to those systems (i.e. models of EU agricultural subsidies). Some provide input into ongoing debates at a more general level which would have impacts on many different policy contexts (i.e. models of the effects of policy on concentration and rate of technological progress). A last group does not specifically talk about specific policy issues at all, but only provides insight into a system which is of interest to policy makers (i.e. models of the effects of Tobin taxes). This plurality of senses of policy modelling makes the task of isolating models to consider difficult. An assortment of models which could serve as further examples in this broad sense of policy agent-based modelling include: Windrum and Garcia-Goni (2008) on the effects of the introduction of ambulatory surgery on an ocular health system; Schwoon et al (2008) on alternative strategies for energy system transitions; Pahl-Wostl & Hare (2004) on water management in Zurich; and Alam et al (2007) on the socioeconomic impact of AIDS in a region of South Africa. The problem with trying to compare all of these approaches is that they address very different phenomena through diverse disciplinary methods. In order to demonstrate the issue more fully, we turn to a specific domain which has been modelled differently by many groups.

The domain chosen for illustrative purposes is the exercise of market power in electricity markets. This example is chosen for a number of reasons. First, it is a market in which substantial changes have been made in the past decades with a shift towards private

delivery. This has had unexpected and sometimes unfavourable results, best exemplified by the California blackouts of 2001 and the blackouts in Eastern North America of 2003. Complexity of structure and unforeseen results are the hallmark of the type of problems which might be tackled by agent-based modelling. Second, it is a domain in which there is both strong incentive and ample possibilities for policy intervention at some level. In terms of incentive, either sharp rises in price or disruption of service represent major issues in which a clear imperative to action exists for governments. In terms of possibilities for intervention, electricity markets typically have regulatory bodies which can make binding decisions. The final reason for choosing this area is that there is the possibility of making a fairly direct comparison between agent-based models and traditional models. A strong tradition of equilibrium modelling of market power in electricity markets exists in the literature. More importantly, it is an area of focus for agent-based models in policy domains. As a very new approach, there are few of these areas which have been approached with this level of rigour (for an overview of existing models, see Tesfatsion (2010)).

The comparison of approaches will be two-fold. First, the types of computational approaches used will be compared. This will lead into a comparison of what types of policy relevant results the models can generate. This comparison is not intended to be comprehensive, either in its coverage of the literature or in its conclusions; rather it is intended only to illuminate how agent-based modelling might diverge from traditional modelling approaches.

Ventosa et al (2005) survey the literature for trends in the modelling of electricity markets. They identify two main classes of models at the market level: equilibrium and simulation models. The majority of the articles which are discussed are in the equilibrium tradition (even in terms of simulation models, many are simply computerized equilibrium models). Consider Borenstein et al (1995), who propose models based upon Cournot equilibrium as generating measures of market power. These are claimed to be more reliable than existing measures, such as the Hirschman-Herfindahl Index. Several observations can be made of the modelling techniques used. They consist of quite theoretical explorations of Nash equilibria. Even within this already highly stylized view of markets there is a bias towards simplifying assumptions for the sake of tractability.¹⁴ For the given, highly stylized scenarios, such models deliver a definite measure of market power. Policy advice also tends to be quite straightforward proposals which aim to directly modify the variables considered. For example, in the Bornstein paper, market power is seen as the result of inelasticity of demand, lack of competitors, and markets isolated by capacity limitations. Thus, solutions focus on increasing the price sensitivity of demand, encouraging entrance of competitors, and building transmission lines to consolidate markets. The question becomes to what extent a policy maker considering a particular policy issue in a particular system might think such a model a good representation of the problem at hand. Also, given the one factor focus on market efficiency, there is little consideration of what effects these

¹⁴ Though the Supply Function Equilibrium (SFE) variant is more structurally accurate, the Cournot equilibrium models are preferred due to their analytic simplicity (the method based on algebraic solutions is chosen over that based on partial differential equations).

changes might have on other criteria for the system (e.g. stability of supply, cost to end consumers).

The smallest step from these models towards agent-based modelling approaches is the slight modification of equilibrium models to make assumptions which increase accuracy of description. On the level of a small step from traditional modelling, Day & Bunn (2001) address the issue of divestiture of electricity generation in the England and Wales (E&W) pool. They focus on the second round of divestiture in 1999, after an initial round was seen as a failure to reign in market power. This led to policy questions of how much divestiture was required in order to mitigate market power. Day and Bunn approach the problem with an approach similar to previous equilibrium based approaches but allowing inclusion of modified assumptions, such as discontinuous supply functions, non-linear elasticity, and asymmetric ownership structures. These changes were argued to more accurately reflect market conditions within the specific market under study. The question posed of the model was: what is the effect on market power of moving from three to five market players? Though this model may be argued to be more realistic than equilibrium models, it remains tied to the concepts developed in this model structure. To the extent it maintains such structures, it remains open to the criticisms levelled against equilibrium models.

A further step towards descriptive accuracy is seen in Micola et al (2008), who explore the effects of reward interdependence in strategic business units of vertically integrated firms. The model is intended to explore integration between wholesale natural gas/wholesale

electricity and wholesale electricity/retail electricity markets. A great deal of prior research is brought to bear on establishing the problem to be addressed and specifying structural components. Studies are cited showing that vertical integration exists in many markets and is being increasingly accepted by regulators. The debate over the advantages and disadvantages of integration from the firm perspective are explored. This is framed via the literature on the most common argument about how market power is exercised: the foreclosure argument. Essentially, this mechanism is based on the refusal of an integrated firm to sell into the upstream market. This will raise the price of the input, disadvantaging downstream non-integrated rivals. For example, if an integrated electricity wholesaler-retailer withdraws from selling on the wholesale market, the wholesale rate will rise. Costs will rise for non-integrated retail rivals while staying the same for the integrated firm, leading to competitive advantage. This argument is criticized by Micola et al. as not being relevant to most actual power markets given that selling to rivals in wholesale markets is often compulsory and that the standard market mechanism of uniform price auction does not allow differentiation between internal and external clients. Instead, the authors look to a more structurally accurate case: strategic business units (SBUs) linked by incentives with netback pricing. Netback pricing is a mechanism in which prices are determined by following the value chain up from retail consumption rather than down from resource prices. This mechanism, which the authors claim is supported by the energy markets literature, is based upon the nature of the particular market: retail prices are stable over longer periods than prices further upstream.

The second literature addressed is that of corporate incentives. The most important conclusion of this literature is that reward interdependence is most effective when tasks are interdependent. Netback pricing sets up large interdependencies between SBUs which implies that reward interdependence will be effective. It is theorized that as a side effect of this behaviour, prices will rise throughout the market. This is caused in the same manner as, but different direction from, the foreclosure mechanism. Given that the retail market gives “price input” to the wholesale market, the retail arm can act less competitively and inflate prices to the benefit of the wholesale SBU. Elucidation of this mechanism is stated to be the paper’s main economic policy contribution.

Despite extensive recourse to the literature to define structural conditions, the underlying model is still quite stylized. It is described as two sequential, multiple unit, compulsory, uniform price auctions representing a wholesale and a retail market. This introduces several simplifying assumptions such as homogenous firm actors, constant marginal costs, a lack of transmission constraints, etc. Specific algorithms are sometimes introduced only in terms of how they simplify the modelling process and sometimes broadly supported by plausibility arguments (e.g. there is a maximum reasonable end-user price, decisions are based on reinforcement learning). Setting of the parameters of the simulation is on a similar basis. Notably, while there is still passing reference to equilibrium concept (i.e. that the model “presents a manifold of non-Pareto ranked Nash equilibria”), these are pro forma and not explanatory. The explanation consists of stepping through the dynamics of the model through the aid of time series graphs of variables.

In terms of policy advice, the model results are used to suggest that regulators who are dealing with integrated SBUs should monitor the bidding behaviour of downstream competitors. This is in contrast to models based upon the foreclosure mechanism, which suggest monitoring of upstream competitors. The developed result is argued to be superior because it is specific to a certain market structure. A further caveat not seen in analytical modelling is that it is explicitly stated that the proposed mechanism does not preclude that a) no such activity is occurring or b) collusion might be taking place on other levels. It is based upon an algorithmic rather than analytical explanation, which renders it more comprehensible but less definitive.

A final step towards accuracy of description is Sun and Tesfatsion (2007) who present a detailed model of the Wholesale Power Market Platform (WPMP), a market design by the US Federal Regulatory Commission which has been implemented in several regions of North America. The modelling effort focuses specifically on accuracy of representation. For example, to ensure that the electrical grid is rendered realistically there is significant recourse to literature on the physical constraints of such a system. In terms of human agents, it is recognized that self-report of behaviour is unreliable given the incentives involved, so observations of strategic behaviour are extracted from human-subject experiments. The market protocols implemented are complex and non-generic; instead of a standard economic class of auction, the actual auction structure of this system is explored. The realistic auction structure is that an Independent System Operator (ISO) takes cost and capacity figures from generators and pairs them with the projected demands of retailers.

This is done in a two part process. Bidding is done on a day-ahead basis, with pricing based upon projected demand. This is supplemented by a secondary process in real time to correct for variations from the projected demand. Using the information at its disposal, the ISO creates optimal solutions (using a location dependant pricing scheme).

This model represents a shift to quite substantial levels of the characteristics of agent-based models highlighted above. It focuses on market performance and incentives to strategic reporting behaviour to market makers within the confines of a realistically rendered structure. Rather than relying solely on generic results in the literature, the model is based on detailed research into specific markets. This level of detail leads to a model with over 30 exogenously set parameters. It utilizes a multi criteria system for judging the market design: efficiency, profitability, sustainable. Computational experiments begin by running the model without strategic behaviour to generate a baseline. This is then modified to explore documented market power mechanisms such as strategic pricing, capacity withholding, and induced transmission congestion.

Given the level of detail in the model, it is unsurprising that its results are difficult to interpret. Even in the baseline case of no strategic reporting, price is determined by the complex interplay of daily load profiles, transmission congestion, and production limits. Introducing strategic behaviour, it is shown that over time, generators learn to implicitly collude in over-reporting marginal costs. Under these conditions, costs are higher but also smoother than the baseline case. This discounts price spikes as a detection method for this

form of behaviour (as had been suggested elsewhere). This suggests that a more direct form of auditing of costs is necessary in order to limit such behaviour. In terms of more direct applications to policy, it is suggested that such a modelling platform might be used for policy purposes through exploration of the performance of actual or proposed market designs from a social welfare point of view (short term reliability and long term investment incentives). It is not clear what the outcomes of such an exploration would be or how they would be communicated to policy makers.

As a brief aside, an indirect result of the Sun and Tesfatsion model might be motivating further research into how to more directly detect the exercise of market power. For example, Entriken and Wan (2005) explore the proposals for a policy measure called the Automatic Mitigation Procedure (AMP), which was proposed by the California ISO to mitigate market power. This is done by altering generator bids which have been identified as being driven by market power as opposed to driven by scarcity. Two examples of such behaviour is bidding in a way which vary with output that is inconsistent with known technical performance characteristics of generation technology and bidding which varies in time without underlying performance or supplier causes. They examine a multiscreen system which examines the conduct of individual actors, the impact of a bid on price levels, and a minimum price at which screening comes into effect. This might be a sort of more direct auditing congruent with Sun and Tesfatsion's conclusions.

Having examined several example models, this exercise has demonstrated the effect of shifting from analytical to agent-based models in terms of policy implications. As models increasingly aimed for descriptive accuracy, the claimed characteristics of agent-based models – the use of boundedly rational, heterogeneous, non-trivially networked agents – became more prominent. There is a clear difference between models in an analytic/equilibrium tradition and those agent-based models which aimed to be highly descriptively accurate. However, the range of assumptions which could be seen within the domain of simulation methods highlights the large degree of variety in models which can be classified under the agent-based modelling tradition. Some efforts were attempts to incrementally improve equilibrium models, some represented attempts which focussed on structural features at a market level, and some represented attempts to replicate markets in a detailed fashion. We can think of these approaches as becoming more and more constrained by a specific context as opposed to theoretical considerations. In terms of policy, these models provide advice at very different levels of specificity.¹⁵ They also represent very different strengths. In focusing on optimal policies, the analytical approach provides strong imperatives to action. Conversely, agent-based approaches seem to provide a stronger relationship to the contexts they propose to intervene in but much more uncertainty as to the strength of advice they can provide. These arguments lead us towards a need to evaluate which models can be considered good representations of the system they are attempting to model.

¹⁵ A further consideration is the extent to which the models used are specifically policy targeted as opposed to simply in policy relevant domains. To be effective, do models which yield policy insight need to be oriented to this task from the beginning?

Descriptive Modelling and Policy

While the idea of making more empirically grounded assumptions is central to current thought on the benefits of agent-based modelling for policy, exactly how model development relates to empirical evidence remains controversial. Some approaches aim to incrementally increase the plausibility of assumptions, while maintaining conventional model development approaches. Other approaches argue that including more empirical content requires fundamental changes in how models are developed. This is seen in the debate among agent-based modellers as to whether models should be driven by a heuristic of simplicity or a heuristic of descriptivity. The simplicity heuristic (the Keep It Simple Stupid, or KISS, approach) was first proposed by Axelrod (1997). Driving this view is a desire to preserve the clarity of causal relationships. With simple models, it is possible to gain a detailed understanding of the observed dynamics. Here, the impact of empirical evidence is restricted to deciding between different sets of simple assumptions. The counter-argument to the simplicity approach is provided by Edmonds and Moss (2005), whose descriptivity heuristic (the Keep It Descriptive Stupid, or KIDS, approach) claims that accurately describing actual systems should be the goal of modelling. The descriptivity claim can be couched in terms of validation. A focus on maintaining descriptivity ensures that a comparison can be made between target system and model to evaluate validity. This is emphasized by the differentiation between easy complexity and hard complexity by Edmonds (2009). Here easy complexity is the production of complex patterns from simple assumptions, while hard complexity is understanding actual observed

complex outcomes. The key argument presented is that at some point, to extend beyond academic interest, models need to address observed complexity.

The descriptive modellers have gained some traction in the debate when it is applied to policy contexts. They argue that because their models are subject to stronger validity tests, they are more relevant to policy (see Moss, 2008). In utilizing the argument that relevance to policy flows from validity, there are connections evident in their approach. In fact, sometimes the descriptive approach is explicitly referred to as evidence-based modelling. There are two points of this argument which are not entirely clear. The first point is that although validity is an important concept, there is not a clearly defined common understanding of validity. Given the importance ascribed to validity, a fuller account of validity seems important. The second point, is the extent to which relevance to policy flows from validity. This connection is not clearly defined, as with evidence-based policy which proposes very little in the way of a mechanism by which models might influence policy decisions. There is an implicit assumption that sound science is sufficient to ensure relevance. We proceed to take these two issues up in turn in the following chapters.

CHAPTER 3: VALIDATION OF AGENT-BASED MODELS

To conclude Chapter 2, there was a discussion of the importance of validation of agent-based models to their use in policy contexts. The topic of validation of agent-based models is of quite recent prominence. Recall that Heath et al. (2009) showed a large change in the importance of validation in the past decade. Despite an increase in interest, no compact literature on the validation of agent-based models. That which does exist suffers from a lack of extensive cross-citation, inconsistent use of terminology, and a diversity of background assumptions about the role of validation in the larger practice of modelling. Different positions on validation are often couched in the disciplinary standards of their modellers, which imply different epistemological standards. The ambiguity of terminological use is further obscured by the often unexplored differences in the aims of different modelling traditions. These patterns are seen, for example, in the Fagiolo et al. (2007) special issue on validation of agent-based models in evolutionary economics. These are specific accounts of how to carry out validation which are not based upon a common agreement of what validation means.

The ambiguity of terminology and the variety of underlying assumptions make it difficult to analyze validity on the level of specific tests. Commonly cited tests of validity in evolutionary economics, such as those discussed in Windrum et al. (2007), are quite open to such difficulties. In particular, two assumptions in these approaches remain unjustified. First, the assumptions that empirical data only comes in the form time series data on

emergent outcomes, with underlying generating processes being unobservable, seems odd. If this were the case, why should agent-based modelling be preferred to analytical modelling in the first place? Second, the definition of validation as reducing the parameter space which needs to be explored only makes sense if we accept the unobservable nature of the underlying processes. Otherwise, it is not clear why parameters alone in the model need to be justified. These issues and others led a descriptive modeller (Moss, 2008) to critique the approach as not being relevant to descriptive modelling. Rather than attempting to discuss validation from the positions of the detailed methods of either Windrum et al. (2007) or Moss (2008), a more fundamental discussion will be entertained. This discussion will not evaluate validity at the level of specific tests, but instead will be concerned with highlighting prominent views on validation and the open debates in the area. Drawing on validation literature from across the social sciences, an attempt is made to develop a framework in which claims about the strength of classes of validity tests can be evaluated at a high level. Specifically, this framework aims to evaluate the claims of openness to stronger tests of validity made of descriptive models.

The first step will be a review of the literature on the definition of validation and its role in the modelling process. A major goal of this will be to align the terminology of different authors. This will lead to a more subtle working definition of validation, which considers different types of validity, validity tests at different stages of modelling, relative degrees of validity, and validity to purpose. With this definition in hand, we will proceed to discuss a taxonomy of agent-based models based upon the relative mixture of theoretical and

empirical elements. On a continuum from quite theoretical models (abstract models) to quite empirically based models (descriptive models), it will be argued that more descriptive models are in fact subject to stronger validity claims.

Validation and its Role in the Modelling Process

At first glance, evaluating validity seems simple enough. A model is valid to the extent that it accurately represents the target system which it is intended to represent. If a model is accurate, it is a reliable representation which can form the basis of action in the world. Upon examination of this simple explanation, several issues emerge. First, there is the issue of representation. A common observation in modelling (see Epstein, 2006; Gilbert, 2008; Nelson & Winter 1982) is that trying to perfectly represent phenomena in a model is a fruitless task. On information theoretic grounds, we can say that any perfect representation of a target system will need to be at least as complex as the target system. Thus, such a project simply replicates rather than explains. Accepting imperfect representation means that choices need to be made as to what is important to represent and this will change the nature of the representation. This selectivity implies biases in the knowledge developed. A first step in clarifying the literature on validation is to differentiate validation from concepts which are related but distinct. First we take up two concepts, verification and calibration, often confused with validation in that they are related to checking the results of modelling. Then we take up a concept, calibration, which is often confused in that it involves the relationship between model and target system. After

clarifying each of these through direct comparison to validation, the fit of the concepts into accounts of the model development process may be discussed.

The conceptual differentiation between verification and validation is well established in the literature. These are commonly confused because they often accompany each other in modelling practice. Verification has been defined as: ~~the~~ “the process of checking that a program does what it was planned to do” (Gilbert & Troitzsch, 2005, p. 22); ~~checking that~~ “the representation is faithful to the simulator’s intentions” (Edmonds, 2003, p. 108); the correctness of model construction or ~~building the system right~~ “building the system right” (Manson, 2002, p. 63). Compare these definitions to those of validation: ~~determining~~ “whether the simulation is a good model of the target” (Gilbert & Troitzsch, 2005, p. 23); ~~determining whether~~ “the expression of the simulation in terms of outcomes is faithful to the relevant social phenomena.” (Edmonds, 2003, p. 108); the truthfulness of the model with respect to its problem domain or ~~building the right system~~ “building the right system” (Manson, 2002, p. 63). As with much of the literature in this area, there are some inconsistencies in these definitions. For example, David (2009) points out that the Gilbert & Troitzsch definition of validation refers to the whole model while the Edmonds definition only to model outcomes.

Despite any differences, what the definitions share in common is that verification involves a relationship between modeller expectations (itself an unclear construct) and model performance, while validation links the model and the target system. Verification is a model-internal process which involves only the modeller and the model, validation requires

external inputs. The key question of verification is: do inputs (I) lead to outputs (O) in the expected manner? The key question of validation is: is the model system (M) a good representation of the target system (T)? A verified model allows legitimate claims about how the model is working; a validated model, about the model being a good representation of the target system. We might expect in practice that these two concepts are related, insofar as the expectations of the modeller are related to knowledge of the target system, but we may have one property without the other; a validated system may be unverified and a verified system may be unvalidated. In the former case, the model is representative of the target but we do not understand why this is so (we may think of neural networks here). In the later case, the model is consistent with how the modeller expects it to work but it cannot make claims of representation (we may think of game theoretic work here).

A second conceptual differentiation is between validation and generation, which sits in a relationship similar to that of validation and verification. However, in implementations the concepts of validation and generation are often mixed, leading to confusion. The concept of generation has become important in agent-based modelling recently. It has been proposed, notably by Epstein (2006), that agent-based modelling utilizes a novel sort of explanation, supposedly distinct from deduction and induction. Generation is defined as the production of macro level regularities from micro level assumptions. The key heuristic question is –How could the decentralized local interactions of heterogeneous autonomous

agents generate the given regularity?”(ibid, p. 5)¹⁶ It is relevant to note the word “could” in the heuristic, as the process is oriented to the production of potential explanations. Indeed, David (2009, p.123), characterizes generation as much like the process of abduction which is the generation of hypothetical explanations to a given explanandum. Grune-Yanoff (2009), argues that generation only produces potential explanations because it simply produces conclusions from assumptions, whose empirical status needs to be established independently. So, while the generative approach is referred to as a “new tool for empirical research” (ibid, p. 4), it is not clear that an empirical connection is inherent in generation.

Generation is related to verification in that it is an internal property of the model. It is a confirmation that the expectations of the modeller (i.e. that certain micro specifications will lead to certain macro regularities) have been met. As with verification, these expectations may be in fact conditioned by empirical factors. On its own, generation guarantees neither the micro nor macro validity of the model. It may be argued that generation implicitly includes some notions of validation. In the full explanation of the generative scheme, Epstein suggests factors which go beyond pure generation. He suggests that the elements to be generated are to be drawn from empirical research and that some variety of statistics may inform evaluation of generative sufficiency. Furthermore, if multiple models are equally able to generate a given phenomenon then they should be judged on the realism of their micro level specifications. What is not made clear is how the generative target is isolated or how comparisons of the micro specification to empirical agent characteristics

¹⁶ If we were to add boundedly rational to this list, the assumptions of decentralized decision making, local interactions, heterogeneous agent characteristics would be just those used in Chapter 2.

are to occur. So, while Epstein's scheme is suggestive of being compatible with validation, the focus of generation is on possibility proofs and not claims about target systems.

Differentiating the concepts of generation and validation, while holding that they are not exclusive in practice, avoids confusion about the empirical status of generative claims.

The final terminological confusion is arguably the most subtle. It is the distinction between validation and calibration (sometimes also called estimation). Calibration is ~~the~~ the process of choosing the values of the parameters that maximize the accordance of the model's behaviour (somehow measured) with the real-world system." (Richiardi et al., 2006, para. 4.19) It involves taking a given model structure and, from some empirical basis, choosing values of parameters. As such, calibration is the process of establishing a relationship between a particular set of model parameters and some body of empirical evidence. A well calibrated model is one where parameter values are empirically justified. Compare this to the given definitions of validation above. While validation relates the model as a whole to a target system, calibration is focused on justification of specific model parameters.

Calibration can be seen as part of the process of validation but it takes the model structure for granted and is only focussed on the initial conditions of the simulation. Any model structure, even one which is completely arbitrary and empirically nonsensical, can be calibrated. This distinction is important, as in the analytical modelling tradition, the model structure is largely set by methodological conventions. This means that validation is nearly synonymous with calibration. The disconnect between calibration and validation can manifest in several different ways. It is possible to calibrate parameters which have no

impact in the target system. Alternatively, a single parameter may be calibrated which is in fact multiple distinct aspects in the target system. Given that the data which is used for calibration (itself often containing unstated biases) is interpreted through the lens of model structure, the use of empirical data does not necessarily imply an empirically meaningful calibration. Conversely, all models have some initial conditions which need to be specified. Some form of calibration might be seen as essential to making strong claims about representativeness of specific outcomes.

A consideration in calibration is whether data is used to directly specify parameters or whether empirically supported arguments are a necessary intermediary. We will refer to the former as direct calibration and the later as indirect calibration.¹⁷ Direct calibration has the advantage of being relatively straightforward. By directly comparing analogues in the model and target system, fairly precise and easily evaluated ranges of values can be generated. But it is the very simplicity of the direct comparison which is also the source of critiques of this approach. Is measured GDP really the same thing as aggregate productivity in a model? To the extent that the model structure does not parallel the structure of the target system, we may question the meaningfulness of such direct comparisons. The other alternative is to make weaker assumptions of representation and attempt to justify ranges of parameters on empirical grounds. There are questions as to whether this type of argument simply hides problems with model structure rather than

¹⁷ Note that an approach called indirect calibration has been proposed elsewhere. The existing approach uses a validation procedure to infer backwards to parameter values (for details see Windrum et al., 2007). This is not the sense intended here.

solves them. One interpretation is that the failure of direct calibration means that there remain errors in the specification of model structure. The issue of how directly data interacts with model structure is a prominent issue in validation as well, as will be explored below.

Having isolated validation from related concepts, we now examine how both validation and related concepts fit into typical modelling practice. A first account of this is by Manson, who suggests the process of modelling is as follows:

Most modeling efforts have underlying theory that is distilled into a *conceptual model*. The modeler then uses *calibration data* to instantiate the conceptual model in a *software model*, in particular, a *multi-agent system model* that represents a *target system*. The multi-agent system model is then subject to *verification*, which involves testing the software model to ensure the proper functioning of its underlying programming. The multi-agent system model is then run in order to create *model outcomes*. *Validation* comes in two varieties, *structural validation*, or how well the software model represents the conceptual model, and *outcome validation*, how successfully model outcomes characterize the target system.

(Manson, 2002, p. 63-64)

In this account, the sense of verification closely parallels that suggested above. In terms of calibration, the accounts are broadly similar. A minor point of clarification is that saying

calibration data is used to “instantiate the conceptual model in the software model” is unclear. Using standard definitions of instantiation from computer science, it is better to say that calibration data is used to instantiate the simulation model.¹⁸ The distinction between conceptual and software model has not yet been discussed. Neither has the differentiation between structural and outcome validation.

A second source on standard practice is Gilbert & Troitzsch's (2005, pp. 18-19) textbook on simulation. They suggest that standard practice begins with identifying a problem to be solved. This leads to definition of the target system which is to be modelled. At this point, observations are made of the system to yield parameters and initial conditions (i.e. calibration). Next assumptions are made and the model designed. The model is then verified to ensure it is free of bugs (i.e. verification). Then validation takes place by comparing model outcomes to target system outcomes. Finally, sensitivity analysis is performed to ensure the model is robust to changes in parameters. This differs somewhat from Mason's view of the process. First, the stages are performed in a different order. Second, there is no explicit mention of a conceptual model. This, it may be argued, is implicit in the stages of choosing the problem and defining the target system. Finally, sensitivity analysis is added to the list of steps. This is the process by which model parameters are systematically explored to determine the robustness of outcomes to specific model settings. As one element of the exploration of model outcomes, sensitivity analysis may be considered a part of the computational experiments stage.

¹⁸ Instantiation being typically defined as taking an abstract blueprint and creating an instance with its own unique variable set (i.e. taking a class and making an object).

Taking these two sources together, we may say the stages of model development are: conceptual model development, software model development, verification, calibration, computational experiments (including sensitivity analysis), analysis, and validation. It is important to consider this explanation as outlining stages to be completed and not as a linear model of best practice. For example, in many software engineering paradigms software model development and verification will be part of a single, iterative process where small pieces of the model are verified as they are written. Two main issues arise from this analysis. First, in Manson's account, both a conceptual and a software model are considered. How are these models related and what does this have to do with validation? The second issue is the different sub-types of validation which are discussed. What is the importance of differentiating output and structural validation? These two issues form the basis of the next two sections.

Multiplying Intermediaries

As noted above, some authors prefer to consider validation as confronting software models with data while others consider intermediary conceptual models. Introduction of the concept of multiple models makes the modelling process less simple conceptually but is arguably more reflective of modelling practice. The process of moving through a number of different models which mediate the connection between target system and software model provides an explanation of the difficulty of generating strong methodological

guidelines. It also suggests the concept of stepwise validation which is a promising conceptual advance.

The idea of step-wise validity has been considered in agent-based modelling before. For example, Richiardi et al. (2006) propose adopting Stanislaw's (1986) three types of validity: theory validity, model validity, program validity. This sets up a chain of validity tests at the translations from target system to theory to model to program. At each point of translation, a judgement of validity may be made. Theory validity examines the link between target system and theory, model validity the link between theory and model, and program validity the link between model and program. While this is an interesting approach, the specific form of intermediation proposed is based upon the rather controversial linear view of the relationship between models, theory, and empirics. So the important point taken is the consideration the quality of representation at each stage of intermediation and the affordances of this to clearer discussion of the complex validation problem.

A less problematic account is provided by Drogoul et al. (2003), who suggest that the modelling process contains three distinct roles (which might but do not necessarily map to different individuals). These roles are: the thematitican, the modeller, and the computer scientist. Each role is seen as leading to the production of a different type of model. The thematician is an expert in the area where modelling is to occur. They create the domain model (in our terms, the conceptual model), which is a narrative account of the dynamics

of real-world agents. This model is taken up by the modeller, whose role it is to formalize the model. This involves understanding both the domain model and the constraints which representation as a finite computer program imply. The modeller is expected to produce a design model, something similar to the software design concepts of a class diagram (a static overview of agents, variables, parameters, available behaviours) and a process diagram (a dynamic overview of how behaviours will be triggered). This model is more formal than the domain model but less detailed than the software model. Finally, the design model is implemented as a computational model in a programming language by a computer scientist. This division of roles is helpful in that it makes clear the different models at different stages and highlights chains of transformations of models. It also suggests a clearer division of labour. For example, a domain expert who has limited computational background will find the domain model much easier to review than if given the final software model. Also, this approach makes verification easier by more formally expressing the ‘intent of the modeller’ in a semi-formal design model. We might view this as replacing Stanislaw’s program-model-theory with the more specific terminology of computational (software) - design - domain (conceptual) models. This provides a similar effect with less a priorism about the roles of theory and data.

A further addition of models is proposed by David (2009). He suggests that in a simulation, three models are involved: a pre-computerized conceptual model, a simulation model and a post-computerized conceptual model. The relevance of this is that simulation is not only mediated in its contact with the target system in the design phase but also in the

interpretation phase. The model acts as a source of consistent and demonstrable ideas for modifications to the conceptual model. A similar, though less analytically precise, structure is suggested by Edmonds (2003).

Consideration of intermediary models makes the path from target system to software model much more indirect. Depending on our interest, we can see the process of validation as being stretched into many distinct phases. Each intermediary model introduces abstractions which may be sources of error and requires separate consideration of validity. In this situation, the task of validation is subdivided into subtasks which might be of more manageable proportions. It is important to note the stepwise approach is not a panacea. Localized judgements of validity are likely to be non-transitive, meaning that even if all individual transitions are validated, overall validity is not guaranteed. The intermediary approach does explain how agent-based models which lack simulation components (c.f. Windrum & García-Goñi, 2008) are partial steps on the way to a software model. More importantly, it provides the possibility of understanding model development in a grounded way.

Types of validation

The second issue that derives from the discussion of standard modelling practice is the varying characterizations of the sub-types of validation. Gilbert and Troitzsch took validation to be a single process, while Manson subdivided it into structural and output

validation. The division between outcome and structural validation is similar to the distinction drawn by Heath et al. (2009) in Chapter 2 between operational and conceptual validity. While the subdivision of validity is common in accounts of validation, the correspondence between categories used is non-trivial. A first resource on types of validation is Tesfatsion's (Tsfatsion, 2009), web-based literature review on model validation and verification. She suggests some of the open questions in this area. Key among these for our purposes is the different kinds of validation which are possible. Specifically, it is suggested that validation may be: descriptive output validation (—matching computationally generated output against already-acquired system data”), predictive output validation (—matching computationally generated output against yet-to-be-acquired system data”), input validation (—ensuring that structural conditions, institutional arrangements, and behavioral dispositions incorporated into the model capture salient aspects of the actual system”), or the iterative participatory modeling approach (a stakeholder driven process involving four steps: —field study and data analysis; role-playing games; agent-based model development and implementation; and computational experiments”). Here, output validation is very similar to Manson's outcome validation. Input validation is similar to Manson's structural validation, though Tesfatsion's definition does not explicitly mention an intermediate conceptual model (as Manson's does). Also, Manson takes model structure as a single entity, while Tesfatsion divides in into structure, behaviour, and institutions. The iterative participatory modeling approach involves multiple rounds of both input and output validation through stakeholder input.

Another similar structure is given by Troitzsch (2004), who takes up Zeigler's (1976) division of validity: replicative validity, predictive validity, structural validity. In terms discussed thus far, replicative and predictive validity map to descriptive output and predictive output validity of Tesfatsion respectively. Structural validity is again more ambiguous, being similar to Manson's structural validity but significantly stronger. It is that the model —not only reproduces the observed system behaviour, but truly reflects the way in which the real system operates to produce this behaviour.”(Troitzsch, 2004) Of note in this definition is the requirement that reproduction of system behaviour is achieved before undertaking structural validation. Also, the requirement of truly reflecting the true operation of the system (as opposed to adequately representing the conceptual model) is quite strong. In discussing these types of validation, Troitzsch notes that structural validation is much more common in the social simulation case. This is attributed to a) a dearth of data due to cost, b) issues with reflexivity of measurement (i.e. measurement of agent attributes can change the attributes), and c) the structural instability (relative to non-living systems) of many interesting social situations. It is also noted that the idea that structural validation occurs after predictive validation is suspect. We have little reason to think this particular order to be preferred, especially given questions about the possibility of predictive validation in agent-based models. It is reasonable then to remain agnostic as to the order in which validation of different types is performed.

A final similar division of validity into outcome and structural components is Yucel and van Daalen (2009). The equivalent of outcome validation is instead termed behavioural

validation. Unlike the approaches mentioned thus far, they identify two sub-types of outcome validation measures (proximity in their terms): numeric proximity and dynamic pattern proximity.¹⁹ Numeric proximity is simply comparison of model time series to target system time series. This is what seems to be implied by most outcome validation. Dynamic pattern proximity steps away from this direct comparison and instead looks at the broad patterns of behaviours. The importance of introducing the idea of dynamic pattern comparison is similar to the discussion of direct vs. indirect calibration discussed above. It raises the question of whether outcome comparison is a direct comparison or one mediated by some analysis of sufficient pattern similarity. In terms of structural validation, there is again another definition. It is seen as —the comparison of interactions and behaviour rules against the existing knowledge about the system being modeled. This knowledge is a combination of theory, empirical observations or the tacit knowledge of the experts.” (ibid, para. 3.17) The role of structural validation is suggested to be situations where the purpose of modelling is understanding processes and mechanisms which lead to behaviours and not just replicating behaviours.

Aside from Yucel and van Daalen’s observations on direct vs. indirect comparisons, there is a broad level of agreement on the concept of output validation. Structural validation is more problematic. It has been defined as: comparison to a conceptual model; ensuring model structures should capture salient aspects of the actual system; reflecting the way in

¹⁹ Three types of proximity are actually identified but the final form, terminal pattern proximity, is a test of whether a certain behaviour of interest has emerged. This is similar to the test of generation discussed above and thus should not be considered validation proper.

which the real system operates; and comparison with existing knowledge of the system, including theory, observation, and expert opinion. In order to justify the importance of structural validation, we now turn to an example showing the explanatory weakness of output validation alone.

The Limits of Output Validation

The insufficiency of output validation alone to support explanation is well demonstrated by Troitzsch (2004). Using the historical case of German school teacher gender desegregation, he provided an interesting example of the divergence between replication and explanation. The model considered was one which attempt to replicate the historical statistics through some very simple assumptions. Relevant statistics show that that the dynamic to be explained was two distinct school populations dominated by male and female teachers respectively which gave way to schools of mixed teacher gender over a timeframe of approximately 20 years. The simple model replicated the pattern effectively on the basis of a process which replaced outgoing teachers with a teacher whose gender was equally likely to be male or female. However, when examining the actual process of teacher gender desegregation through historical methods, the process was found to be largely driven by a merging of schools, which was in turn driven by economic factors. As the percentage of girls in grammar school increased, small towns in the region under study could not afford to support separate schools for boys and girls. Through the process of merging schools, staff was merged leading to mixing of male and female teacher

populations. The key point to take from this example is that even though descriptive output validation was performed, the model provided a poor explanation of the underlying process. Agent mergers (in this case, mergers of schools) were not included in the model structure, thus no amount of calibration or output validation would improve the explanation. If we imagine the process being extended to future data sets, supposing these data sets are drawn from populations with similar underlying mechanisms, we would expect the model to be predicatively validated but again a poor explanation. This brings into question the sufficiency of these types of validation for system intervention. Trying to use the mechanisms described in the model as the basis for policy would be highly misleading.

Here we see the weakness of output validation to support algorithmic explanation. Given the flexibility of agent-based modelling, it becomes rather easy to modify assumptions iteratively through feedback from output matching. This provides very little assurance of any similarity between the process in the model and the process in the target system. Criticisms of only replicating and not explaining have been levelled against, for example, history-friendly modelling.²⁰ A second observation is the explanatory weakness of plausibility assumptions. The problem, as discussed in the section on generation, is that many different plausible assumptions can be made. Worse, it is not clear that discussions of plausibility get us any closer to actuality. Following Edmond's (2009) argument from

²⁰ The appropriateness of history-friendly modelling is related to its representative claims. If it only claims to match the outcomes of historical processes then the criticism of not explaining is apt. However, if it is seen as an attempt to represent structures through evidence gathered by historical investigation, then the explanatory case is much stronger.

Chapter 1, we might agree that plausibility tests allows models to engage at the easy complexity level (showing how simple agents can interact to create complex outcomes) but not at the hard complexity level (showing the actual mechanism which led to an observed complex outcome).

Validity as Relative and Purpose Dependant

As validity is a matter of checking the fidelity of representation between a target system and a model system, it may be judged in degrees. Such a concept of relative strength of validity, as opposed to a dichotomous variable, is a source of some confusion in the validation literature. David (2009) points out that this is partly due to the fact that validation and verification have very different meanings in the philosophy of science and the practice of modelling. Traditional usage in philosophy relates verification to the problem of truth of models and validation to internal consistency and a lack of known or detectable flaws. In the sense used here, questions of validation are closer to the philosophical concept of confirmation, which is simply consistency with available evidence. This terminological clarification is helpful in reminding us that a valid model by our definition does not make claims about truth, only about consistency with available evidence.

The relativistic concept of consistency allows consideration of the question of Tesfatsion (2009): —Must the intended purpose of the model be known before meaningful empirical

validation can proceed?” In a more declarative sense, Forrester takes the position (in discussing the validation of system dynamics models) that “validity, as an abstract concept divorced from purpose, has no useful meaning.” (Forrester in Sterman, 1984, p. 51) Some philosophical definitions of models accept this as well. For example, Maki defines the modelling process as: “Model M is an entity used by agent A to represent target system S for purpose P. The inclusion of purpose or function suggests pragmatic constraints on the required respects and degrees of resemblance. The desired sort of resemblance is a function of the uses to which models are put, the purposes they are supposed to serve.” (Maki, 2005, p. 305) The idea of validity to purpose is stronger than the observation of Heath et al. in Chapter 2, that propensity to validate relates to purpose. It claims that the nature of validation also changes with purpose. This implies that different models may be required for different purposes for the same target system. A further implication of the consistency conception of validity is that it is an ongoing process of building up confidence in the use of a model for a particular purpose. When introducing a model, researchers may well be expected to make claims about validity. However, such claims are only a first step in assessing validity. It is only through continuous testing and use that confidence is built up. Assessment of validity is not a one time proposition but a partial and ongoing process.

A Working Definition of Validation

To this point, this chapter has reviewed the literature surrounding the validation of agent-based models. Though this is a fast moving field and likely subject to rapid change, this

review at least served two purposes. First, it outlined issues which any consideration of validation must contend with. This included recommendations to make a clear distinction in theory (if not in practice) between validation and related concepts of verification, calibration, and generation. Further, it requires specification of whether validation includes the model structure or only the outputs of the model. Models employing only output validation are vulnerable to critiques of only replicating and not explaining. Structural validation is a controversial topic with open questions of how to evaluate structural similarity, including specific tests used and how strict the similarity must be. One possible method of evaluation is the stepwise approach which suggests judging the preservation of structure through a series of intermediate models. Finally, it is asserted that judgements of validity are relative and centred on consistency with available data rather than truth. This suggests that rather than a single model of a given system, there may be many different models for different purposes.

This leads to a discussion that is more instrumental to the discussion in Chapter 2. We may now propose a working definition of validation for purposes of policy modelling.

Validation is the process through which a model is evaluated to determine the extent to which it is an adequate representation, for its stated purposes, of some target system. It may test adequacy of representation on the basis of ability to generate certain outcomes and/or on the basis of structural similarity between the model and the target system. When considering structural forms of validation, it may be helpful to adopt a step-wise form of validity testing. By examining transformations of the model on the path to formalization

into a software model, deformations in structural conditions may be identified. It is also relevant to provide working definitions for related concepts. Calibration is the setting of parameters of the model from an empirical basis, taking the model structure for granted. It may be either direct (taking target and model parameters as literal analogues) or indirect (using empirical support but not making direct target-model comparisons). As opposed to validation and calibration which relate to a target system, verification and generation are taken as model-internal processes. Verification is the process of ensuring a bug-free translation of a conceptual model into a design model and then into a software model. Generation is taken as the process of confirming that micro specifications lead to macro dynamics in the software model, in the manner specified in the conceptual model. If a model has been verified and generation confirmed, then we have a consistent and well-understood potential explanation. These definitions significantly clarify the concept of validation.

Heterogeneity of Agent-Based Models

In the preceding discussion, the meaning of agent-based model was held constant for the purposes of exploring the literature on validation. Having established a working definition of validation, it is pertinent to discuss the extent to which agent-based models can be taken as a homogenous category. There are many different directions which explorations of model heterogeneity could take. Given the importance of the debate between simplicity and descriptivity in Chapter 2, the motivating heuristic of modelling is taken as an

important differentiator. However, this implies a rather fragmented framework. In order to support a more general framework, we turn to a recent approach in the philosophy of modelling.

Models as Mediators

Attempting to define systematic relationships between agent-based models and empirics is a difficult task. This may be partially due to the contentiousness of the concept of model in the philosophy of science. The debate about the relationship between models, theory, and empirics has a long pedigree. It is not our purpose to engage with this debate here.

Instead, we will introduce in passing a recent current in the literature on modelling which is closer to our concerns. The concept of models as mediators is outlined in Morgan and Morrison (1999). What differentiates this approach from many other philosophical approaches is that it attempts to begin the philosophical debate much closer to observed scientific usage. This leads to concepts and terminology which are much more compatible with that of model practitioners than many philosophical discussions. It allows us to develop a definition of models in the context of how they related to theory and empirics.

The definition of models as mediators is developed from several detailed case studies. The concept of mediator is non-trivial and turns out to be to have similar themes to Actor Network Theory (ANT) discussions of technology. Models are seen as “autonomous agents and instrument of investigation.” (ibid, p. 10) In terms of construction, models are

seen as being made up of disparate elements such as theoretical, empirical, mathematical, and analogical. Models are seen as being inherently underdetermined with respect to any type of element.²¹ Thus, even the most theoretical models cannot be constructed by theory alone. Similarly, models which are ‘data-driven’ or phenomenological are partially independent of the data which underlies them. Given that process of model building is first and foremost the assembly of bits, it is the place in which mediation between different domains occurs.

Framing models as mediators does not provide a strong normative theory of modelling practice. It suggests that relationships are complex and not fully grasped in previous ways of thinking about models. Though not directly specifying a way forward, the theory does provide some heuristics for moving forward. First, to attempt to catalogue the diverse elements that contribute to model construction in specific cases (i.e. to examine modelling practice carefully). Focus on constitutive elements, implies a less abstract form of categorization, which is open to revision or extension on the basis of ambiguous or anomalous cases. Second, given that Model-Theory-Empirics relationships differ across different modelling traditions, it is reasonable to hypothesize that validation approaches and results will also vary. Debates which neglect this factor and attempt to speak to agent-based model as a single class of models are likely to engage in methodological debates which are irresolvable at the methodological level.

²¹ Note the ANT undertones of underdetermination and agency through assemblage.

A Taxonomy of Agent-based Models

The import of this for modelling practice will be shown by its compatibility with existing agent-based modelling taxonomy. Using this taxonomy, we will discuss how validation might proceed in each case and the implications of this for the types of knowledge produced. The established view of validation will help to clarify what types of validity are being tested and the extent to which these types of validation are appropriate to particular types of models.

It has been argued that it is useful to approach the categorization of models upon their constituent elements. An examination of the literature on modelling practice shows evidence of use of similar concepts by modelling practitioners. One example of differentiating validation practices according to model type is provided by Gilbert (2008), a manual on the practices surrounding the building of agent based models.²² He argues for defining models at three levels: abstract, mid level, and facsimile. Abstract models are presented as tools for theory development, with little direct relationship to empirical contexts. Mid level models represent particular social phenomena but in a sufficiently general way that it may be applied widely. Facsimile models focus on describing contexts as accurately as possible. The word facsimile seems to be somewhat pejorative, so it is perhaps more neutral terminology to refer to them as descriptive models. This taxonomy

²² It is relevant to note the author's centrality in the field of social ABM (for example, he is the editor of the *Journal of Artificial Societies and Social Simulation*). He suggests that validation is a difficult and controversial area in which conclusive research is yet to be forthcoming. This tends to confirm the contention that validation is an area of recent interest but little agreement in the social modelling community.

of models is partly based upon Troitzsch's (2004) argument for three different types of models: those which represent an abstract class of systems, those which represent the qualitative behaviours of instances of systems, and those which represent the specific behaviours of instances of systems. Having established the taxonomy, we can now discuss their relationships to validation

Abstract models, according to Gilbert, have the primary purpose of advancement of theory. Therefore, these models should be evaluated in a similar framework to other theory generation mechanisms. Gilbert envisions this framework as providing only a very weak form of validation.²³ The only real link to empirical contexts is a judgement of the plausibility of the agents involved.²⁴ Rather than attempting to further validate within the realm of theoretical models, Gilbert sees the input of this plausibility as instrumental to development of successful middle range theories. In this framework, strong validation is never applied to theoretical models themselves. Such models should be seen as sources of consistent and theoretically sound hypotheses but not reliable knowledge.

In terms of the validation process discussed here, abstract models are often satisfied with a hypothetical conceptual model. In seeking similarities between disparate domains, they often take large transformative steps without validation. For example, some evolutionary models consider firm and organism evolution in the same model structure. Focussing on

²³ Some of the validation steps proposed do not fall within our definition of validation at all. These are model-internal processes of checking generation and sensitivity analysis.

²⁴ This is just the tie-breaking mechanism of Epstein (2006); if two models equally generate a given phenomena, the one with a micro specification most in line with knowledge about agents should prevail.

similarities and ignoring differences allows for theoretical advance but it also substantially weakens prospects for validation. One weak path to introduction of empirical information is the use of so-called stylized facts. These are intended to represent the important elements which are to be explained. Heine et al. (2005) have framed stylized facts as not replacing empirical testing but as heuristic to avoid modelling for the sake of modelling. We could think of stylized facts as the mechanism to the hypothetical conceptual model. Here output validation is a weak test. Structural validation is very weak, as it is generally accepted that the structure is not representative. The advantage of this approach is that verification is strong and the model mechanisms well understood.²⁵

Turning to descriptive models, a much clearer connection to validation is possible. These models are ~~intended~~ to provide a reproduction of some specific target phenomena as exactly as possible...” (Gilbert, p. 43). Two different streams of descriptive models exist. The first aims at making plausible predictions. An alternative use of descriptive models is the descriptive modellers described in Chapter 2. These modellers eschew predictive modelling and instead take descriptivity as a goal in of itself.

Using descriptive models for predictive purposes is largely a matter of ensuring output validity. Comparison of the target system data and model data is complicated by the fact that model data is typically a set of data traces from multiple runs, with different initial conditions, parameter settings, and/or random seeds. Conversely, target system data is

²⁵ For the affordances of highly theoretical models, see Frigg (2003).

often a single historical trace. In non-linear models, small variations in parameters and initial conditions can lead to very different results. Even with the same settings, simply changing the random number seed can greatly affect the outcome of stochastic processes. Rather than attempting to choose any individual data trace, a distributional understanding is necessary.²⁶ Even with a distributional understanding, the so-called problem of historical uniqueness remains. Briefly, this problem is that if we get a result which is divergent from the target system history, it is unclear if this is due to a poor model specification or of the actual event being improbable within its possibility space. Gilbert sums up the situation well:

If it [the most common outcome] does not [match actual observations], one might wonder whether this is because the particular combination of random events that occurred in the real world is an outlier and, if it were possible to ‘rerun’ the real world several times, the most common outcome would more closely resemble the outcome seen in the model! (ibid, p. 44)

If the goal of the exercise is only related to outputs, it is difficult to justify using agent-based models over any other type of model.

The non-predictive stream of descriptive modelling is more interesting conceptually as it introduces element of structural validation to the discussion For example, Moss (2008)

²⁶ Multiple runs with different {initial conditions, parameters, random seed} sets must be done and comparison of the outcomes performed.

claims that this approach is focussed on just those practices which are germane to evaluation of validity claims. Iterative comparisons between model and target system allow for inconsistencies with the target system to be identified. These comparisons are typically done on both an output and a structural basis. The use of this anomaly generation mechanism, allows for a working definition of structural similarity. A model is similar to a target system if an informed observer cannot generate any elements in which they disagree. Note the relationships of this method to anthropological method more generally. Note again that the strongest this approach can say is that the model is consistent with available observation. It may fail to capture elements because of incompleteness of observation, lack of conceptual framework for searching for anomalies, or structural change of the system under study. Given the need of a well defined target for comparison, these types of models have typically examined phenomena like Natural Resource Management in geographically limited regions. Though the definition of the structural validation steps leaves many elements open, it does demonstrate the prospects for validation of models that claim rather direct forms of representation.

It is relevant to note that in both forms of descriptive modelling, verification tends to be weaker than more abstract approaches. This may be due to the interest of the researchers, the rather complex structures of the model not being amenable to testing, and the typical skill sets of researchers involved in these projects.

The final type of model to consider is middle range models. Treatment of these models is the most tenuous of the three. This is likely due to the lack of a strong heuristic which identifies the approach unlike the abstract and descriptive approaches. A middle range model is meant to “describe the characteristics of a particular social phenomenon, but in a sufficiently general way that their conclusions can be applied widely to, for example, most industrial districts rather than just one.” (Gilbert, p. 42) In validation of this type of model, we are most interested in qualitative resemblance. It is suggested that given the lack of intent to directly represent particular instances, it is not meaningful to do detailed output validation. Given this, he proposes some methods of characterising qualitative resemblance. First, we might compare the dynamics of the model to real world dynamics. Does the “history” of the model seem meaningful in relation to empirically derived histories? A second form of qualitative similarity is a comparison of the statistical signatures of the model and target. For example, if two properties are known to exist in a power law relationship, this should be borne out in the model. While Gilbert uses the term qualitative resemblance, this terminology is somewhat confusing. In the first sense (model history as meaningful relative to target system history), it is a very inexact term. In the second test (of statistical signatures), the test is in fact a form of quantitative analysis. It seems that what is meant is something more akin to the differentiation between coarse vs. fine grained agreement. Using qualitative descriptions or quantitative data, the goal of validation at this level is coarse grained agreement. In the terms of Agar (2005), we might think of this as more-or-less validity; as verifying that we have captured the differences that make a difference.

Mid level models are difficult to discuss with any degree of precision. Given the mix of empirical and theoretical, we would expect that such models would lie somewhere between abstract and descriptive in terms of affordances. That is, we expect they would be stronger on validation than abstract models but weaker than descriptive models; similarly, it is reasonable to expect they are weaker on verification than abstract models but stronger than descriptive models. We might expect their conceptual model to be at least somewhat validated with respect to their target. Similarly, we expect at least some test of step-wise validity as this model is formalized. Whether this necessarily implies iterative validation is unclear.

The key result of dividing models in this manner is an isolation of the affordances of different types of models. Descriptive models are open to stronger validity tests because they claim representativeness. Abstract models offer more clearly understood mechanisms because of their focus upon parsimony and generality. Given the approach of segmenting upon empirical/theoretical inputs, this result should not be surprising. It is a recapitulation of a well worn debate between theoreticians and empiricists. Rather than engaging in this debate, this framing is an attempt to emphasize the different affordances in terms of validation of theoretical and empirical traditions. The strong validity affordances of descriptive models form the starting point of Chapter 4.

It is notable that the framing of the validation affordances of models presented here is significantly less sanguine than the claims made of models in the evolutionary economics

literature. Models are often assumed to have undergone validation on the basis of justification of plausibility of assumptions, some form of calibration, and output validation. For example, the Windrum et al. (2007) review paper is in this tradition. The analysis here suggests that use of such tests alone may be appropriate to abstract modelling, but that they do not achieve the level of validity which is often implied. In order to make stronger validity claims, it is not only improved tests of correspondence of data models which are needed. Rather, a key step is a more focused definition of the target of the model. The idea that prospects for validation are related to the strength of the representative claims being made is a neglected area of research.

Conclusions

This chapter has discussed the definition of validation and the extent to which validation potential varies with model orientation. After discussion of the debates surrounding validation, a working definition was offered. This was: the process through which a model is evaluated to determine the extent to which it is an adequate representation, for its stated purposes, of some target system. An important element of judging adequacy of representation was the difference between outcome validation (the replication of certain outcome patterns) and structural validation (a judgement of structural similarity between the model and the target system). Implicit in this definition is recognition that validation can only evaluate consistency with available evidence. The degree and respect of consistency varies with the purpose of the model. One element of purpose is suggested to

be the abstract vs. descriptive role of the model. A taxonomy of abstract, mid level, and descriptive models were considered. Based upon its strong claims of representation, it was concluded that the descriptivist claim of being open to stronger validity tests, presented in Chapter 2, is well founded. The implications of this for the use of agent-based models in policy contexts form the subject of the next chapter. Before turning to this topic, it is relevant to note some important conclusions relating to evaluation of validity in of itself.

One important point is that validation will nearly always imply making do with imperfect data sets. This may spark additional empirical research to test more directly the claims of models. As Manson notes, “The most cogent result of validation may be illustrating the need for more validation data since this need can define how further experiments should be conducted.” (2002, p. 70) Thus, validation is not only a process of assessing models but also a driver of the direction of empirical research.

Another issue which deserves highlighting is the extent to which validation needs to extend beyond outputs. An attachment to output validation as the dominant form of relation to empirical contexts is deeply entrenched due to philosophical attachment to the relationship between prediction and explanation. As such, it enters many validation schemes as an assumption that output validation is the only relevant type of validity to consider. Given the flexibility of agent-based modelling to generate outputs, we should be very surprised if there are not several different ways of generating the same output. This is an issue which accounts of validation in evolutionary economics have not fully come to terms with. It is a

crucial issue as it goes to the heart of algorithmic explanation and the difference which adopting the agent-based approach makes.

CHAPTER 4: VALIDITY AND RELEVANCE TO POLICY

Chapter 2 ended with the assertion of descriptive agent-based modellers that a focus on validation is essential to improve the relevance of agent-based models to policy. In attempting to define validation in Chapter 3, it was found that judgements of validity are non-trivial. However, it was argued that we can classify agent-based models by the extent to which their content is intended to be constrained by phenomena. Where intent to faithfully represent is strong (i.e. at the descriptive end of the spectrum) there is the possibility of strong validity tests. This can be seen as being aligned with evidence-based policy, as both see validity as a variable to be maximized. The best model for a given policy task is just the model which is most validated. In this case, the path towards improved policy models is clear: large investments should be put into the fieldwork and related descriptive modelling processes.²⁷ We should go forth and build up a patchwork of domain specific models which, taken together, would provide a comprehensive picture in which to ground decisions. However, the focus on validity as the primary criterion of policy relevance remains unjustified. The case for relevance is almost entirely on the back of rejection of highly theoretical (abstract) models. This negative heuristic in modelling, like evidence-based policy generally, does not provide a mechanism of how such models will be useful to improving policy decisions. While it is granted that models in the highly abstract, optimal policy tradition are problematic, the lack of positive arguments is also

²⁷ Note that though this argument is similar to that of Brown et al in Chapter 2, it is suggesting very different models. Despite the framing of Brown et al as making more realistic assumptions, the descriptive approach would reject these models as ungrounded and lacking connection to a particular target system.

problematic. Here we will suggest that models developed in the descriptive tradition are limited in their relevance to policy through the very source of their validity. Exactly by trying to represent specific systems, they deliver knowledge which is not directly useful in making decisions.

The Local Tool Critique of Descriptivity

The argument for the descriptive approach to agent-based policy modelling relies heavily upon making a comparison with abstract modelling. In order to create a positive justification, we might examine the practice of descriptive modelling for policy. This is not possible as, despite the rhetoric, there are few clear examples of models which claim to be directly useful to policy makers. It will be suggested here that this is due to elements inherent to the method. In examining the domains in which descriptive modelling has been applied, there has been a strong bias towards use in local contexts. The most developed usage is in the context of natural resource management, especially in areas which experience renewable resource failures due to a lack of coordination over common resources. Beyond such contexts, there is little positive evidence of successful descriptive modelling. Following the descriptive approach requires a great deal of fieldwork, notably extended and intensive interaction with target system elements. This is self identified by descriptive modellers as a barrier to adoption of the approach. What is less well addressed is how this direct tie to particular target systems not only presents a temporary barrier to adoption (by requiring more work per model) but also restricts the contexts in which the

method can be applied. What may be termed the local tool critique of the descriptive approach is that it is inherently restricted to rather narrowly defined domains. This would in turn provide an explanation of the lack of obvious path of application to policy contexts. If the policy relevant target system is seen to be outside of what can be modelled within the descriptive approach, then it is unsurprising that the application is not straightforward. Obviously, this would call into question the argument for descriptive modelling in policy more generally. To explore the local tool critique, we turn first to a more in depth overview of the practice of descriptive modelling.

The priority of descriptive modelling in social contexts is generally given to the French school of companion modeling (ComMod). ComMod utilizes an iterative three step model which directly integrates stakeholders (c.f. Barrateau et al (2003)). These steps are: fieldwork to characterize the target system, development of appropriate models, and computational experiments on the model. The computational results lead to new questions for fieldwork, reinitiating the cycle. In this approach, participants directly interact with models and modellers to inform continuous model development. The model structure is directly specified and calibrated from a particular group of subjects. As mentioned in Chapter 3, this direct interaction allows for strong validity claims. What is underemphasized in discussing this approach is the extent to which a relatively circumscribed problem domain is required to produce detailed representations of particular instances. These constraints have led to the ComMod approach primarily focusing on exploring quite geographically localized issues such as the D'Aquino et al. (2003) model of

a Senegalese river valley and Becu et al. (2003) model of a small river catchments in Thailand.

These observations are borne out by examining another example of this approach, the Etienne et al. (2003) model of the land use patterns of different groups in the Causse Méjan plateau in Cévennes national park, France. The central controversy was the preservation of a rare grassland ecosystem in the face of different forms of land use. Representations of the vegetation structure of the area were derived from empirical surveys. Key stakeholder groups were identified from studies of existing land use. Agent characteristics were defined directly from a survey of individuals on their relevant behaviours, resources, and attitudes. The database resulting from the survey was used to directly initialize runs of the models. Models were developed separately for each group of stakeholders, leading to negotiation on the basis of multiple alternate scenarios. This eventually led to the adaptation of a common model framework. The methodology utilized relies on intensive interaction with stakeholders, making it difficult to see how the methodology could be used in a wider geographic area, for example to address land use for the whole of France.

If descriptive modelling is restricted to local contexts in practice, this puts into serious question the extent to which it can be applied to many policy areas. The perception that descriptive models are only applicable to rural and/or institutionally static situations is a major issue if we are interested in influencing policy in a highly complex and interrelated societies. While there has not been concerted work responding to this critique, there has

been an implicit response from descriptive modellers which suggest that this critique can be overcome. For example, Moss (2008, para. 1.6) argues that this localization is not inherent in the descriptive approach. This is supported by citing models of issues of larger scope, such as water management in southern England (Downing et al, 2000) and Afghan power structures (Geller and Moss, 2007). While these models are clearly influenced by the ComMod tradition, this response to the critique fails to address divergences of the example models in terms of their directness of representation. Qualitative data is certainly still important in the definition of the models but is often introduced only after some analysis. This is not entirely clear in the examples cited above, as they do not provide a detailed account of the development process and data inputs.²⁸ We now turn to an example which does discuss the integration of data explicitly which demonstrates the shift of methodology that addressing target systems of larger scope entails.

Alam et al. (2007) describe the impact of HIV/AIDS on the socioeconomic structure of the Sekhukhune district of South Africa. The model consists of interleaving networks at the individual and household level. Agent and mechanism design is intended to be constrained by independent evidence, following the evidence-driven (descriptive) agent-based modelling approach. However, several intervening factors make direct representation of the situation quite complex. The authors note that the stresses of HIV/AIDS are difficult to separate from other stresses such as: food insecurity, climate variability, market

²⁸ The Afghan example cited by Moss was actually (Geller, 2006), which is a German language publication. Perhaps the development process is stated more explicitly here but is not analyzed here due to the insufficiency of the author's German.

fluctuations, and levels of government and non-government (remittance) support. Detailed data on the interaction of all of these factors is not available. Instead, the evidence sources utilized are the expert opinions of an external case study team and secondary literature (which provides a mixture of qualitative and quantitative data). This is a clear case of introducing mediating qualitative analysis to bridge the gap between available data and the level of detail necessary to constitute the model structure.

Examining the model structure in further detail shows that the model is not determined directly for data, as in Etienne (2003), but from the conclusions of the analysis of data. Structural features of the model (i.e. the definition of household by cohabitation as opposed to biological relationship and the process by which orphaned children join existing households) are based upon the analysis of qualitative research. Some elements of the model are represented by distributions derived from the conclusions of secondary literature rather than direct characteristics (i.e. the process which decides who becomes infected with HIV/AIDS in each period). Parameters such as the proportion of infected mothers who pass the disease on to their children are similarly defined. At the boundaries of the simulation, significant breaks from direct representation occur. For example, agent health is presented abstractly as a proportion of full health. Another example is the representation of community structure by a small world network of low clustering coefficient (a plausible but certainly not directly specified step). These examples are identified as areas for future research, constrained by a dearth of data.

While available evidence is still the focus of this model, the amount of information which would facilitate direct definition decreases. Incorporating a larger scope has changed the available evidence. For example, it is claimed that “Modelling complete household economy for our case-study region is not possible as this requires meticulous fieldwork.” (ibid, para. 5.15) This sheds some light on the claim of the paper that agent behaviours are “built on descriptions of specific and identifiable actors”. (ibid, para. 1.1) The qualifier ‘built on’ disguises the difference in methodology which is occurring. The representation of qualitative conclusions, along with detailed calibration, is different from the more direct ComMod approach. In the ComMod approach the actors are specific descriptions; in the expanded approach, actors are representative of important classes of agents (themselves based on specific descriptions/expert opinion). This is not a critique of the Alam et al. approach, but rather a more careful analysis of the extent to which the practice fits with the ‘evidence-driven’ or descriptive categoricals. This distinction is blurred by a rhetorical approach which conflates all models which incorporate qualitative empirical research in order to create a contrast with models driven by “prior theoretical consideration or developed ad hoc on the basis of introspection and/or speculation” (ibid, para. 1.2).

The Alam et al. model answers the local tool critique but it does not do so within the methodology of the descriptive approach. The difference of methodology is in the degree to which the models attempted to be representative of particular empirical objects. Representation is the most direct when model elements clearly represent particular empirical objects. In the case of the Etienne et al. model, the human agents were modelled

after particular participants and the environment was modelled after a particular geography and ecosystem. Directness of representation is attenuated as mediating analysis of the particular empirical objects is introduced. Alam et al. moved in this direction by moving away from directly representing individuals. The logical endpoint of this movement is models in which mediation is such as to render the representation untraceable to any particular empirical object, even if it derives from research into particular empirical objects. This may be called indirect representation.²⁹ To explore this, a model is considered which is not traditionally associated with the descriptive approach but which seems to fit the definition of indirect representation.

To anchor the indirect approach in a specific example, we need to ask where the descriptive method might go if applied to problems of an even larger scope and the resultant decrease in available data. A plausible next step in this line of thought is Agar (2005) who models the properties of drug addiction in urban America based upon ethnographic research in the area.³⁰ The model focuses on the role of drug use stories or “buzz” in driving addiction cycles. Two agent characteristics are seen as relevant: risk and attitude. Risk is a parameter of agents which signifies their willingness to acceptance of unknown experiences. On the basis of extensive empirical work in the innovation literature, it is argued that this characteristic is normally distributed. Attitude is a variable which describes the positivity

²⁹ Note that indirect here does not refer to the relationship between researcher and subjects. To develop qualitative insights, researchers are still assumed to interact with particular individuals. It may be the case that the particular qualitative methods used will be of varying intensity of interaction with increased scope, an idea expanded on further in the discussion.

³⁰ As the model does not fit within existing traditions, Agar is not particularly careful about framing. He frames the model as being of drug use dynamics, full stop. In fact, the model is more appropriately framed as being of drug use in urban America because this is the situation from which it is derived.

or negativity of an agent towards a particular illicit drug. This is varied by observing the “buzz” generated by other agents in the environment. Another set of assumptions comes from prospect theory, which argues that avoidance of negative results is preferred to the attainment of positive results. The choice model of the agents is simple: if they are addicted, they consume the drug and if risk exceeds attitude, they consume the drug.

In this model, it is not possible to claim direct representation of any particular drug or social context. Instead the claim put forward is that the representation aims to reflect ‘emic’ differences.³¹ The conclusions underlying the specification of parameters and mechanisms are derived from interaction with subjects, along with broader input from characteristics of opiate addiction. It is a representation of the most important aspects of drug addiction generally and not any drug addiction specifically. It is not possible to empirically justify the modelling choices by direct reference to any particular example. Individual differences and differences in aggregate circumstances are papered over, the parameters being only valid as rough approximations. It is also recognized that operationalizing assumptions were necessary which were beyond the scope of what could be directly justified by the data. This was argued to be done according to a two part rule which implied selection of appropriate non-empirical elements. First, the selection of appropriate theories was directed by important emic differences. Second, theories were

³¹ The distinction between emic and etic understanding in anthropological work roughly corresponds to insider and outsider perspectives, respectively. So a representation from the emic perspective should be meaningful to the people it is intended to represent.

used which have proved valid in many different contexts and therefore can be argued to be fairly universal elements of human behaviour.

This approach can be seen as a further step in introducing indirectness. The Alam et al approach was argued to represent the conclusions of qualitative research rather than the data itself. Agar steps away from representing the detailed qualitative structure to representing only the structures which qualitative research had suggested to be the most important. While this approach still has some resonance with the descriptive tradition in its links to detailed study of specific systems, its claims of representation are significantly different. The mediating research processes transformed the data so that it is disconnected in the final model from any specific instances.

The local tool critique has been developed and explored using example models. Based upon this argument, we may conclude that the descriptive approach, in the strong sense of direct representation, seems quite open to the local tool critique. Only by moving away from the ideal of direct representation was it possible to address larger contexts.

Implications of Directness of Representation

The argument about moving away from direct representation opens up the question of what is meant by ‘larger contexts’. It may be useful to attempt to define more formally what the ‘local’ in the local tool critique is. This will be argued to revolve around notions of scope.

Having this definition in place, it will then be possible to speculate on the relationship between scope and directness of representation. This, in turn, will permit commenting on the broader question of the prospects of the descriptive approach to support policy.

To this point, the term ‘local’ has been used in the intuitive sense of limited size and complexity. It is relevant to introduce the concept of scope as a characterization of the goal of representation. Scope is a difficult term to define exactly. The simplicity tradition of arguing for an understanding of underlying dynamics as opposed to a “detailed account of utterly accidental details” (Bak, 1999, p. 10) might be seen as an argument about scope. The critique of descriptivity as a local tool is certainly about scope. As a first approximation, we could imagine scope being operationalized by a measure like number of subjects participating in the target system. This seems misleading though, as it disregards the level of complexity of participation. Rather than providing an inadequate absolute definition, we instead formalize some of the factors which contribute to relative differences in scope. Some factors which might contribute to increased scope are: number of agents, degree of variety in agent characteristics, complexity of agent aggregations/networks, and degree of dynamicism of agents/behaviours. All of these, of course, suffer from definitional issues but they are useful as heuristics to consider scope. Though a more formal definition of scope is an important future project, this heuristic sense is sufficient for the current purpose.

Note that scope, as defined here, is a choice about setting the boundaries of the target system. It is a claim about the representative goal of the modeller and not about the structure of the target system itself. For this reason, it is chosen in preference of scale, which has connotations of being a property of the target system. Limited scale has implications of the target system being composed of a constrained social or geographical area. It is interesting to note that a critique similar to the local tool critique has in the past been levelled against anthropological method more generally with respect to its restriction to the analysis of village level social phenomena. This may be seen as having been partially overcome, as the method has in fact been applied to complex aspects of life in ‘developed’ nations. The key move was taking the constraint of scale (the village level) and understanding that it is only one way of reducing scope.³² A similar effect of constraint could equally be from selecting a portion of a large body. For example, in examining bureaucratic culture, a manageable scope might be achieved by choose to study a small bureaucracy (small scale) or a specific section of a broader bureaucracy (a portion of a larger scale). The point here is that the direct approach requires an intensity of interaction which restricts the scope.

Having defined the key terms, we can now discuss the importance of the relationship between them. It was implied in last section’s examination of specific models that increasing the scope had implications for directness. Here, an attempt is made to more explicitly argue for this trade-off. The argument can be made on two levels. First, there are

³² Note that limited scale implies limited scope but not the reverse.

practical issues with the level of available data at a given scope and level of resourcing.

This can be understood as a matter of proposing feasible projects. The second issue is more foundational and relates to the desirability of directly specified, increased scope models.

This relates back to arguments about the role of models as tools which intend to provide clear explanations.

The practical argument, or weak hypothesis, for the scope-directness trade-off is on the basis of feasibility of a research project and not absolute constraints. It might be helpful to think in terms of the Project Management Triangle, prominent in the popular press of the project management community (e.g. Newell & Grashina, 2003). The triangle is made up of the three main deliverables of any project: speed, quality, and cost. As the theory goes, a given project can only have two of the three. Any project that is done quickly and well will be expensive, quickly and cheaply will be of poor quality, well and cheaply will take a long time. We are not interested in the details of this approach but rather its focussing power on the essential trade-off which occurs in the pursuit of some end. Its benefit is in a) specifying the impossibility of meeting all criteria simultaneously and b) providing some idea of what sort of effect modifying one variable might have for the others.

Consider a modified version of the same model which focuses on the possibilities of descriptive modelling. The three poles might be: resources, directness, and scope. Here

resources are composed of human resources and time.³³ With sufficient resources we might deliver a model which is directly specified and of large scope. Using the previous example of drug use, if infinite resources were available many different individuals using many different drugs could be studied. But if resources are limited, as they almost always are with such projects, a trade-off must be considered on the directness-scope axis. If scope is relatively low, we can do a fairly direct specification of the agent characteristics. But as we aim to investigate problems of larger scope, we need to rely on empirical data which has already gone through some level of analysis. In this version of the trade-off, we might see indirectness as a necessary evil to begin addressing problems of a certain scope. Through time and investment, it could be incrementally addressed. As a first step in exploring a problem of increased scope, an indirect approach is a much more feasible option.

The other possible barrier to directly specified, high scope models, the strong hypothesis, is more epistemological. Questions remain with such models about their communicative and epistemological properties. They may be open to the simplicity critique of being so complex as to be incomprehensible. These arguments would say that the point of modelling is to understand more clearly and a lack of judgement from the researcher as to what to include in a model simply muddles the situation. This critique is avoided at a low scope but becomes increasingly relevant with increasing scope. Many questions present themselves. Is there some value in having an increased scope model as a baseline for studies of decreased scope? How then, would one model be explained by another model? These

³³ This may be proxied by number of researcher hours. It could be many researchers in a short time or a few working through time.

questions and many others await further study. We do not address the epistemological argument in detail here, as this is a complex question. However, it is notable that an effort to more formally ground the concept of scope utilized here seems a prerequisite to considering such an evaluation.

To summarize, the scope-directness trade-off may be characterized by a weak/practical hypothesis (indirectly specified models as a practical intermediate step) or the strong/epistemological hypothesis (indirectly specified models as the most effective tools for increased scope). Regardless of which hypothesis we accept, for problems of increased scope, an indirect approach is a reasonable first step in model development.

The Case for Mid Level Policy Models

If we return to the taxonomy of introduced in Chapter 3, the shift towards indirect representation can be expected to result in a movement from a descriptive approach and towards a mid level approach. In fact, the notion of progression from direct to indirect representation provides some clarity to the ambiguous concept of mid level models. Before discussing the case for mid level models, we first return to evaluate the definitions of mid level modelling previously discussed.

Impact on Defining Mid Level

Recall that in Chapter 3, the taxonomy of models was largely based on Troitzsch (2004) and Gilbert (2008). Troitzsch (2004) argued for three different types of models: those which represent an abstract class of systems, those which represent the qualitative behaviours of instances of systems, and those which represent the specific behaviours of instances of systems. The indirect method suggested here focuses on the important behaviours derived from particular systems, which is the second of Troitzsch's categories. Gilbert (2008) argued for defining models at three levels: abstract, mid level, and facsimile. Abstract models are presented as tools for theory development, with little direct relationship to empirical contexts. Mid level models represent particular social phenomena but in a sufficiently general way that it may be applied widely. Facsimile models focus on describing contexts as accurately as possible. In the indirect approach, there is a link to an empirical context but direct description is not claimed. Again, the approach offered tends towards the mid level. These definitions of mid level were criticized for being fairly ambiguous. Mid level can become a conceptual catchall for models which do not fit within the well developed simplicity and descriptivity approaches. The approach here of grounding the mid level in specific examples, makes it clearer how the mid level differs from the descriptive. Troitzsch's concept of representing qualitative behaviours of instances might be rephrased as representing conclusions, derived from the detailed study of particular instances, on the most important mechanisms. To Gilbert's claim of "may be applied widely", we note that representation is not claimed beyond the instances from

which the model is generated. It is certainly tempting to think of such models as hypotheses for investigating what appear to be similar instances, but there is no *a priori* reason to assume that the claim of similarity is more than a hypothesis.

One important clarification of mid level modelling is its relationship to the claim which is sometimes made that evolutionary economics focuses on mid or meso level explanation. In the context of evolutionary economics, mid level describes the scale at which analysis occurs. This positioning is relative to macro economics (which addresses entire economies) and micro economics (which addresses the firm level). Mid level in terms of agent-based modelling is describing the level of connection of model elements to particular target systems. This positioning is relative to abstract approaches (which have weak connections to particular target systems) and descriptive approaches (which have strong connections to particular target systems). While there is some commonality to these definitions, there are also important differences. According to the directness-scope trade-off discussed above, there is good reason to believe that as the unit of analysis grows, there is an accompanying decrease in the level of information available to directly specify the model. Thus, at higher scale it is likely that it will not be very fruitful to attempt a directly specified model. However, there is no such dynamic in terms of abstract modelling, which can be applied at any level. Indeed, in neoclassical economics, modelling is abstract regardless of the level.

The Case Against Directly Specified Descriptive Models

The case against directly specified descriptive models largely relies on the local tool critique. Having seen that the existing defence against this critique was to build models which became more indirect, the critique does not seem to have any substantive response. The scope-directness trade-off provides a theoretical justification for why we might consider the application of directly specified models to high scope problems to be likely. The issue is well stated by Gross and Strand on the prospects of building a model of high scope with realistic microstructure: “Trying to build a model with the purpose of true and valid representation of the system’s microstructure may in many cases be likened to try to get to the moon by climbing a higher tree.” (Gross & Strand, 2000, p. 30) The notion that it is extremely difficult to build directly representative models of high scope may explain the lack of descriptive models addressing non-local contexts. Indeed, the arguments for applying a descriptive approach are largely based upon overcoming perceived shortcomings of existing modelling approaches. There is a lack of positive examples of models which demonstrate this theoretical effectiveness. While this can be argued to be due to descriptive models not getting a fair hearing due to prevalent methodological norms (which is likely to be a part of the explanation) a supplementary explanation exists. This is that the scope of policy models is largely set by the policy problem being addressed. Beyond the weak hypothesis (and its implications for the impractical level of resources necessary to produce the huge number of necessary models), there are issues of the type of knowledge being produced. There is a disconnect between the outcomes of descriptive

modelling and the required inputs for improved decision making. This suggests that there is a significant gap between the level at which we are capable of providing direct representations and the level at which decisions need to be made.

Rethinking Relevance to Policy Contexts

In the introduction to this chapter, the idea of considering validity as the dominant consideration in the evaluation of policy models was raised. This can be seen as essentially asking: is “good (empirical) science” enough to ensure appropriateness of modelling approaches for policy purposes? The missing factor was consideration of relevance of the model to the decision context of policy makers. While relevance is an obvious factor from the perspective of policy analysis, there is little work on its systematic relationship to validity. Even Yucel and van Daalen (2009), who discuss the importance of balancing relevance and validity, offer little analysis of how they are related. They frame relevance as important but obvious factor for experienced policy analysts. While policy analysts may have an implicit sense of relevance, there is value in explicitly relating relevance and validity. The idea of scope developed here has implications for both validity and relevance to decision making contexts. If approaches which focus on validity and approaches which focus on relevance imply different levels of scope, this has important implications for policy modelling.

From the perspective of decision making, it is relevant to ask if models below the level at which decisions are being made are very useful. Returning to abstract models, the critique of which spawned the descriptive modelling movement, it was argued that general notions of how systems work are not terribly useful to make decisions about specific, historically embedded systems. It is relevant to ask the opposite question as well. How useful is it to understand just one particular subsystem of the system in question? Does it help in making technology transfer policy to understand how one particular institutional setting operates? Perhaps more accurately, given the achievable limits of the descriptive approach, does it help to understand how technology transfer within some particular sub-field of some particular institution operates?

The argument for a mid level approach is that descriptive approaches may have some role in advising policy but a rather indirect one. Not unlike abstract models, they do not provide a direct way forward but may offer insights to approaches that do. That an approach may be indirectly useful is a far cry from the claim that it is the way to do policy modelling. It may be argued that in order for the indirect effects to manifest, there must be the direct conduit. This is just the models whose scope is defined by the policy problem at stake, a scope often above that achievable by descriptive modelling. In this view, when creating models which are intended to be directly relevant to policy, scope is not a free variable but a hard constraint which modellers must do their best to work within. If there is anything to this hypothesis, then a focus on development of mid level models which

indirectly specify agent characteristics will be very important to the future of policy modeling.

Case Study: Mid Level Modelling

The case presented for the application of mid level modelling to policy has been methodological. Though the discussion has been supported throughout with reference to example models, the true test of the methodological case will be in the development of models within this tradition. Fortuitously, an ongoing empirical research project at the University of Calgary offers promise of becoming a testbed of this approach. Though the model has not yet been developed, prospects for such an approach look rather positive. We will start the discussion by examining how a mid level model would approach the issue in question compared to alternative approaches. This will lead to a discussion of the primary conclusions of qualitative work completed to-date. This will suggest the outlines of a conceptual model.

The topic which the model looks to address is the dynamics of local innovation systems, with an eye towards implications for economic diversification. How do the dynamics of local innovation relate to moving into new activities and/or industries? Consider how different approaches would proceed through research to policy advice. Abstract modelling would attempt a general model of innovation and explore its effects on some model indicator of diversity. This might be based on a series of stylized facts about innovation

(most innovation as incremental and only a small minority being radical, a normal distribution of propensity to innovate among agents). It would then look to generate some stylized facts about diversity (perhaps some statistical regularity drawn from an industry classification system like the North American Industry Classification System (NAICS)). This approach would be open to questions of validity on several fronts. Another possibility is undertaking descriptive modelling, which might start into the problem by choosing a case study of a diversification event and try to work backwards to the innovative source. Or it might choose some small subset of firms in an area and try to explore their innovative behaviour. Regardless of the specific approach, it would only address some small subset of the problem. It is not clear that this small subset would help policy makers to make diversification policy. Finally we might try mid level modelling. This would approach the problem by first performing research at some scope to characterize the key dynamics of the system in question. Identified dynamics would form the basis of a conceptual model, which could be formalized into software model. A research project at the University of Calgary exists which has tried to explore the dynamics of a particular local innovation system which could form the basis of such a model. We turn to this data now.

The qualitative research which underlies the modelling interest of this chapter is the contribution of the University of Calgary research group of Drs. Langford and Hawkins to a multi-university research network on the dynamics of local innovation systems. The Innovation Systems Research Network's project on social dimensions of innovation in cities (ISRN2) is made up of research groups from 22 Canadian universities. The present

data is based on data from Calgary, a city of approximately 1 million inhabitants, best known as the hub of the Canadian energy industry (e.g. it is home to the head offices of 85% of oil and gas firms). The projects in the network proceeded by exploring three themes: the innovation behaviour of firms, the role of creative individuals (c.f. Florida, 2002), and the role of civic infrastructure (municipal government, cultural, charitable, and civic organization). For an overview of the project, see Wolfe (2009).

Data collection consisted of N=123 semi-structured interviews, conducted between 2006 and 2009. Interview subjects were identified based on selection of a sample representative of the statistically reported employment diversity. This yielded a dataset of wide coverage but limited depth. These interviews were then transcribed and coded by content analysis and tagging of topics raised in the basic innovation questionnaire guidelines of the ISRN2 project. For the current purposes, the most important aspect of the data collection and analysis process is the type of output they provide. They yield a picture of relevant factors and dynamics which is only indirectly connected back to individual actors. The focus is not on providing descriptions of specific instances but of generating structural or mechanism knowledge.

From this process, the clearest mechanism uncovered was the gravitational attraction of the central oil and gas cluster. Despite a fair degree of diversity according to standard employment concentration measures, it was found that innovations identified among creative individuals and firms were strongly focused on support to the central cluster.

Where other clients become important, it is frequently a matter of spin-off from the oil and gas focus. For example previous studies found that the wireless and global positioning systems clusters in Calgary initially grew from requirements of the central cluster (Langford et al, 2003). Several examples were uncovered where innovations drew on the oil and gas related knowledge platform. For example, a diagnostic imaging company at first seems to have little relationship to oil and gas. However, knowledge inputs from geophysical imaging are applicable to diagnostic imaging. This is a good example of the main outcome of the study: successful innovation in the region is related to the knowledge platform of the dominant cluster. As indicated by the examples, firms which utilize knowledge related to the central cluster to innovate for other markets are, on average, more successful in Calgary than those which do not. Moreover, this is likely to be the mechanism by which the diversity of the economy may increase.

The primary conclusion of the qualitative study was that success in innovation was conditioned by knowledge relationships to the central cluster. This does not necessarily imply that the innovative outputs are in the same industry as the central cluster. Examples exist in which a new sector which develops to service the central cluster becomes independent. There is also knowledge transfer to innovations distant in terms of products. Here, we attempt to formalize this taking the Firm as the unit of analysis. Firms rely on some Knowledge Base in order to produce some Product. These might be represented by position in a Knowledge Space and a Product Space, which exist in some non-linear relationship. Thus, clustering in the knowledge space has non-trivial implications for

clustering in the product space.³⁴ Two mechanisms exist for change. First, agents perform a “leverage-capabilities” function which causes them to search in their local knowledge neighbourhood for new opportunities. From our study, this represents the majority of innovative activity. Second, new agents are introduced via the “new-knowledge-base” function. This represents entrants into new areas of knowledge, whether by explicit investment in knowledge creating institutions or chance. The calibration of the model will be the strong knowledge clustering in the knowledge space seen in our data. Running the model as a baseline, we will then be able to experiment with two different policy options for economic diversification. One scenario is New Knowledge Bases, which focuses on building up knowledge in areas distant from the present knowledge cluster. A second scenario is Strategic Knowledge Base Expansion, which focuses on selecting from opportunities around the central cluster on the basis of product diversity opportunities. It is hypothesized that investing in related knowledge will actually lead to more product diversity.

In terms of policy implications, two observations jump out even at this preliminary stage. First, it points out the extreme difficulty of trying to setup in an entirely new area of knowledge space. This points out why the now largely discredited practice of attempting to recreate the Silicon Valley cluster was bound to fail. Second, it points out that not all ventures into the product space are equiprobable to succeed. For example, IT which supports the knowledge structures of the central cluster is much more likely to succeed than

³⁴ Note that traditional industrial concentration analysis focus on product space diversity.

attempts which go far outside current knowledge capabilities. It suggests an approach of surveying current knowledge strengths (both explicit and tacit) and building into related areas. In this case, perhaps environmental engineering or clean technologies are candidates. Interestingly, these sectors may expand even with a diminishment of the originating sector, suggesting that the knowledge platform may drive output in areas which are not strictly correlated to the market cycle of the central knowledge platform. This example highlights the differences between approaches which strategically leverage existing resources and those which attempt to develop every possible industry. Note that the conclusion in favour of strategic diversification would not tell policy makers exactly what to do. It would only suggest that investments in knowledge capacity should consider both product diversification and relationship to the central knowledge platform.

Given such a model and its interpretation, further investigation of the existing data set will be possible. The model will serve to provide new directions in which to interrogate the target system. Beyond the existing data, it suggests thinking further about how to systematically categorize the knowledge and product spaces for different companies.

Would such a multidimensional map bear out the conclusion of a strong central knowledge cluster with a diverse product profile? This suggests that the model has possibilities for moving the empirical project forward. It further suggests a possibility for validating the model. A second possible path to validating the model is to create a process to bring the model back to interview participants for feedback. This will require some thought, as the indirectness of representation suggests that a more indirect form of feedback will be

required as well. Though the question which we would want answered is “is this a good representation of the system?” we might have to ask different questions to get there.

Another intriguing possibility is that it may allow exploration of a broader thesis. This thesis is that the knowledge platform effects seen are characteristic of city regions which have a dominant central cluster. This could also be explored in the opposite direction as the hypothesis that having a dominant central cluster is characteristic of mid-sized Canadian cities. An available contrast is other study sites of a similar size where dominant (but very different) central clusters exist. A discussion of the similarities and differences of the mechanism knowledge derived may be mediated by the model. This would provide interesting results on the importance of the domain specificity of central cluster (i.e. is a central cluster which focuses on coordination of natural resource extraction and processing essentially similar to one which focuses on government, information technologies, etc.³⁵

Implications of the Mid Level Shift

It has been suggested that indirect specification of agent-based models from qualitative study is an appropriate approach for policy studies and a more precise definition of mid level models. This shift in focus is likely to have many effects, some of which are explored here.

³⁵ Given the theory that secondary clusters are usually related to the dominant knowledge platform, it is suggestive as a hypothesis that some secondary clusters in Ottawa (perhaps the Information and Communications Technology cluster) exist because of strong relationships to government funding possibilities. This is entirely speculative.

A Way Forward for Incrementalist Policy?

The fundamental assumption of the incrementalist approach is that policy makers, like economic agents, satisfice. In Lindblom's words, they muddle through. The term muddle is perhaps not entirely accurate. They do muddle in the sense of not being guided by a totalizing framework which clearly demarcates the right policies. They do muddle through in the sense of engaging in a simultaneous process of learning about means and ends.

Where the term muddling starts to fail is in its implications of being a random or undirected process. Similar to the debates about the cognition of economic agents, there is a great deal of space between completely random local search and completely structured global search. If muddling was completely Brownian, then models could have little effect on the process.

The satisficing argument in terms of policy should be that the frame of reference for interest in models is set by the decision to be made. In Lindblom's terms, the models can move the debate some distance down the strategic continuum. This is a much more complex dynamic than the approaches which see policy modelling as simply applied science. It is arguable that the evidence-based approach attempts to move from the optimal policy view of applied theoretical science to the view of applied empirical science. But in accepting the framing of applied science, it runs into similar issues.

An open question is what models based upon these principles may provide. Returning to Lindblom's characterization of the policy process, we may hypothesize some modes of

interaction. First, models may provide reinterpretations of the status quo. This may include different evaluations of the effects of current policies, shifting the boundary of what is included and excluded in the marginal analysis, or revision of definitions. Second, models may identify new scenarios for incrementalist analysis. Third, they may provide clearer analysis of the marginal differences between policies. Finally, a point not addressed explicitly by Lindblom, they may provide longer range strategic projects upon which the tactically necessary incrementalist method can act.

The evolutionary approach may also feedback into shaping the mid level analysis. In terms of the example discussed here, the question is sharpened by asking what are the marginal differences between an approach which attempts to leverage a platform related diversity and one which strikes out in all directions.

Beyond Familiarity

One element to discuss is the extent to which the approach presented here directly impacts policy modelling debates. As mentioned above, the introduction to a recent special issue on policy modelling (Dawid and Fagiolo, 2008) took the line that familiarity was important to inducing agreement of policy makers. The thesis, which can be caricatured as “if it looks familiar, policy makers will believe us”, was partially correct. Stated negatively, “if it looks unfamiliar, policy makers will ignore us”, it is probably quite correct. Abstract models with little connection to the problem domain are not likely to help make policy

decisions. However, the idea of familiarity conflates validity and relevance. On the one hand, familiarity implies a level of face validity, which is a prerequisite to getting attention in the first place but insufficient for taking action. Familiarity also implies a structuring of the problem congruent with current scenarios and in line with actionable changes.

Validity effects

In introducing the concept of descriptivity, its strong ties to concerns over validity were noted. Given the methodological modifications, it is relevant to ask how the move to more indirect methods affects claims about validity and reliability. If the ease of evaluating the validity of a model is related to the directness of the comparison (i.e. Moss, 2008), then the becoming more indirect might be seen as having a negative effect on validity claims. Just how negative this effect might be for different uses is an open question. It might be helpful to explore this in the context of moving away from the idea that every research project at a given level will necessarily lead to the same result. As with all qualitative study, there is room for reasonable disagreement (i.e. if the study was replicated by a different team, at a different time, etc. the results would vary). Agar (2004) argues that we can imagine different studies in an abstract space with studies of different but equally valid results defining the boundaries of reasonable disagreement. This idea is complicated by the fact that different models are likely to approach the phenomenon at different levels of directness. We turn to this now.

An interesting question to consider is how models at different levels of directness might be compared. This is similar to, but more complex than, issues of docking models. Attempts to dock different models are attempts to reconcile models which have different internal structures but are intended to model the same thing. If we add the directness issue, the process involves reconciling models which are modelling subtly different things. This may be seen as an issue of encapsulating one model within another. Whether this is appropriately done by integration (including the more directly specified as a submodel of the less directly specified) or summary (using the conclusion of the more directly specified to modify the less directly specified) is another open question. A further open question is how indirectly specified models might interact with more abstract models which do not have an explicit basis in direct observation.

Normalization of Expectations of Policy Makers

Unlike both the optimal policy and evidence-based policy approaches, mid level models do not claim to provide scientific solutions to policy problems. Rather they may best be thought of as providing structural accounts which are somewhat grounded. The role of policy modelling moves away from telling policy makers what to do on a day-to-day basis and towards trying to provide a framework for evaluating policy decisions. In incrementalist terms, it may give the important variables upon which marginal analysis may act. Even here, a given policy model may be interpreted to introduce different

variables for marginal evaluation. For example in the case study explored here, the structure of having one dominant knowledge platform may be seen as being undesirable.

This role of framing rather than solving problems runs against the standard perception of what expert knowledge is able to accomplish. The issue of how to normalize the expectations of policy makers is not an insubstantial one. If this approach is to be taken up by policy makers there is a need to challenge claims that solutions can be provided. This is problematic on two fronts. First, the idea of scientific policy advice is a convenience of political rhetoric. Claiming to base a policy on science shuts down uncomfortable political confrontations. Second, there is an incentive for those giving policy advice to overinflate the weight of their analysis. If a policy maker is faced with several voices saying they have the right answer, it is very difficult to argue that no one could possibly have such definitive answers. Perhaps the best strategy will be to play to an understanding of what policy makers already know: policy decisions are made for political reasons. Framing policy analysis as a framework in which politics can play out may actually be liberating. This is certainly an open issue which requires further thought.

CHAPTER 5: DISCUSSION AND CONCLUSIONS

This paper has attempted to answer the question of whether agent-based modelling is more appropriate than alternative forms of modelling in supporting policy analysis. In order to answer this question, a great deal of ground has been covered. It is relevant to briefly summarize what has been argued.

Chapter 2 provided background and definitions of agent-based modelling, examined theories of academic input to policy, and reviewed the state of the nascent literature on the subject of agent-based policy models. It was argued that existing frameworks for understanding agent-based policy models are lacking in explaining why agent-based models in particular are relevant to policy. Through the example of modelling electricity markets, the implications of shifting from analytical to agent-based models were highlighted. In terms of their inputs to policy, shifting away from ideals of analytical simplicity and towards emergent complexity increased the validity of the models but decreased their ability to generate simple imperatives. The issue of validation was highlighted as an important issue in determining the reliability of models but also as an area of continuing controversy.

This led to validation being the topic of Chapter 3. The literature surrounding validity and its relationship to the process of model development was reviewed. An important conceptual difference was found in the differentiation of outcome and structural validation.

It was proposed that in the absence of structural validation, outcome validation can indicate a model generates similar outcomes without any assurance that it provides a meaningful explanation of dynamics. In attempting to speak to a systemic relationship between types of models and their possibilities of validation, the taxonomy of abstract, mid level, and descriptive models was taken up. It was suggested that descriptive models derive their constraints most strongly from phenomena and thus exhibit the strongest possibilities for validation.

To begin Chapter 4, the validation-centric descriptive approach was seen, like the evidence-based policy approach, to fail to offer an explanation of why scientific validity leads to policy relevance. The local tool critique was proposed as a term to describe the bias of descriptive models towards geographically and/or socially limited models.

Anchoring the discussion in existing models, it was suggested that there exists a trade-off between scope and the directness of representation of a model. The weak version of the trade-off was premised on arguments about the resources necessary to undertake the required work with increasing scope. The strong version of the trade-off raised questions about the contribution to explanation of directly specified, scoped up models. Based upon this trade-off, it was suggested that mid level models represent the possibility of most directly relating to policy contexts. This was explored through the preliminary stages of development of a policy agent-based model on the Calgary innovation system and the prospects of diversification strategies. In terms of future work, it was suggested that there are possible synergies between the mid level modelling approach and ideas of evolutionary

or incrementalist policy. Finally, the implications of the thesis presented here were discussed with reference to exiting debates.

Artificial Policy: Development and Prospects

Applying agent-based modelling to policy continues a tradition of attempting to supplement traditional policy making processes with formal models. These models are fundamentally tools for policy makers, so the normative framework of technology might be a useful parallel. Policy made with the help of such tools may be termed artificial policy. The term artificial has three possible senses. At its worst, policy may be artificial in the sense of a poor imitation of the real thing. This interpretation sees attempts to improve the tool base as a distraction from the essentially discursive task of policy making. At its best, policy may be artificial in the sense of an improved version of the “naturally occurring” version. This interpretation sees improvements in the tool base as the driving force behind improving the process of policy and making it more scientific. The third view is a relativist stance. This view sees previous versions of decision making as just as artificial as the new forms. The new form of mediation may be better or worse than the previous versions, and the best we can do is to compare the different forms of mediation in terms of their marginal differences.

The alternatives against which agent-based modelling support to policy must be compared are those which are discursive and lack models entirely (e.g. the pure incrementalism of

Lindblom, 1959) and those supported by traditional analytic models. Given the critiques of being completely incrementalist, there is reason to believe that approaches which lack model support may fail to see appropriate connections and generally lack direction.

Modelling can help to extend analysis beyond what would have been done in a discursive framework. On the other hand, the analysis presented here tends to indicate that traditional analytic models do not fair well on evaluations of validity. The same can be said of certain forms of agent-based models (i.e. those in the abstract category). Though the practice of agent-based policy modelling is not yet developed to the point of substantially stronger validity claims, these models at least offer the possibility of debate and advancement on the validity front. The same cannot be said, if history is to be any guide, of traditional modelling approaches. By balancing the connections possible through greater focus of formalization with the validity of maintaining at least an indirect link to the systems on which policy is being made, agent-based modelling seems to offer some advantage. This makes agent-based modelling, and especially mid level approaches, an appropriate technology.

Role of Descriptive and Abstract Models

The general argument that has been made is that neither abstract nor descriptive models are not appropriate for policy. Abstract models (e.g. sandpile models of economic evolution such as Bak, 1996) raise questions about validity of representation. Descriptive models (e.g. the village level models of the ComMod tradition) fail to attain sufficient scope to be

useful in making policy decisions. Though the future of policy modelling may not be the sandpile or the village, this does not mean that such models will play no role in policy contexts. Abstract and descriptive models may still be indirect inputs to policy makers as background knowledge. We turn to this now.

Abstract models have been argued to be at too high a level to be generally useful for policy purposes. Contra the Brown et al. approach, simply moving from implausible to plausible assumptions does not move us as far as is necessary along the validity axis. It also suggests that in moving to truly empirically valid models, we move away from the generalist, world economy model approach to economics. ‘_Generalist’ and ‘_valid’ in this context are seen as being incompatible. Though this line of thinking may seem to be divergent from thought in the natural sciences, it comes to a similar conclusion. For example, though in principle all climate dynamics are based on fairly well validated physical processes, understanding these micro dynamics are insufficient to provide useful information on the behaviours of the system. Where can abstract models be of use? Consider the rather abstract model of Chapter 7 of Victor (2008). Here it is shown that the assumption that micro efficiency gains lead to overall decreases in energy use is not as straightforward as it might seem. Both logically and historically, there is evidence to suggest that gains in micro efficiency have ambiguous effects. For policy, these abstract models can suggest areas where assumptions may be suspect or unexpected connections exist which require further research.

The critique of abstract models in policy is not novel. Many in the evidence-based policy community have been making similar critiques for a long time. The contribution of the argument presented here is to suggest that descriptive models do not make wide enough claims to be useful beyond micro domains (e.g. the management of small ecological regions). It does not seem reasonable to expect that sufficiently detailed descriptions will ever be sufficient for policy making. Rather, it suggests that we have need of representations which can be improved by the results of detailed research. If we ever hope to reassemble the social, or, to use a less lofty goal, to leverage the knowledge of researchers to better decisions in the world, we need to move beyond descriptive models. Where can descriptive models play a role? Given the argument that descriptive models are the strongest in terms of validity claims, they may be useful in engaging particular controversies. For example, in the mid level model developed here, there might be controversy around whether a particular industry conforms to the model. This could be the target of a more detailed, more descriptive model. Another interesting area of research will be to use the results of detailed models to improve more indirect models.

Status of output of mid level policy models

While there is a long tradition in policy modelling of attempting to provide exact forecasts, the approaches discussed here focuses on what was termed algorithmic explanation. As discussed in both Edmonds (2003) and Yucel and van Daalen (2009), there is a move towards dynamic structural knowledge. Rather than an exact answer we might expect at

most a constrained list of possible trajectories. The level of constraint provided will allow proposal of mechanisms which offer coarse-grained agreement with empirical data. In their focus on structural as opposed to output elements, a very different kind of advice emerges. This advice is less specific than analytical advice but also much better grounded in specific situations. It is able to enter a dynamic relationship with field work which was quite lacking in analytical models which make no attempt at structural validity.

Conversely, since such models do not aim for exact representation, a level of theoretical assumptions will necessarily be involved. This definition of mid level does not specify what the exact mix of elements will be. It is an open and interesting question as to how to relate the results of models at different level.

Contribution to the evidence-based policy debate

Some of the strongest arguments for evidence-based policy revolve around concepts of explaining observed dynamics of actual social systems. Part of this movement is due to the clearly erroneous advice derived from vastly simplified models. This can be seen as the motivation behind Edmond's differentiation between easy and hard complexity. In attempting to be more reliable, the modelling process is pushed towards validation-centric frameworks. However the case for evidence-based policy is largely negative; it essentially argues that if unvalidated models led to bad decisions then more validation means better decisions. The failure of this could be termed the policy modelling paradox: unvalidated models may lead to poor decisions but increasingly validated models do not necessarily

lead to increasingly better decisions. There is a complex dynamic between validity and relevance. The localization factors which ensure the validity of descriptive models also make them less relevant to practical decision making. Rather than evidence-based, perhaps the best we can do with policy is evidence informed. It seems unlikely that the weight of evidence will ever be sufficient to preclude human decision making. Science will not make normative choices for us. It also reminds us of Nelson and Winter's remarks that negative advice will usually be stronger than positive advice. Furthermore it reaffirms the importance of work on participatory methods. If models do not lead in the direction of scientific decision making, then social decision mechanisms are necessary to complement analysis. Better modelling does not, as some more optimistic proponents such as Brown et al., mean the end of ideology. Perhaps the best we can hope for is curbing of the most erroneous elements of ideology.

Future Work

This thesis has considered the role of agent-based modelling in influencing the advice given to policy makers. A novel aspect of this account is that it was developed from an explicit view of the specifics of the modelling technique and of the policy making context. As a preliminary attempt to fuse these areas of study, it did not include substantial work on the details of implementation. A first element of future work is to discuss in greater detail how existing modelling methodologies can be modified to move towards the mid level. This may include further developing the details of the outline presented in Chapter 4 as to how

descriptive approaches can be adapted to attain greater scope. Additionally, it may take the form of asking how abstract approaches can become more empirically based. This will likely take the form of considering abstract models as hypotheses about specific systems, by making stronger representative claims about more focussed target systems. With the generation of such models, it will become necessary to give a fuller account of both the validity and relevance of mid level models. In terms of validity, there are questions of how to evaluate claims about capturing the most important mechanisms. In terms of relevance to policy, the hypothesis that models should not be seen as directly giving specifying the answers to policy questions deserves further work. This will focus on a more extensive theory on the structuring influence models can have on policy decisions.

Prognosis

Returning to the original question, agent-based modeling could be a very powerful tool for supporting policy making. However, concerns about validity suggest that research supporting policy should draw upon study of specific contexts. Despite this, taking explicit account of the policy process has led to the conclusion that descriptive modelling is not the best way forward for policy modelling. Due to its inherent locality, descriptive modelling in the direct sense does not deliver results of sufficient scope to be amenable to direct interaction with policy. These two requirements, of validity and of relevance to policy contexts, need to be balanced. Given this tension between policy relevance and validity, it has been suggested that middle level models represent an underexploited niche for policy

studies. By enforcing some checks on validity, the grossest errors of scholastic and ideological delusion may be avoided. By examining models which proceed beyond specific instances, a mechanism to input to policy is opened up.

The fundamental thesis here is that an approach which attempts to attack the problem at the same level that the decision maker faces it is more likely to be a useful tool. This approach is fraught with dangers which parallel those of decision makers. There is never enough information to be sure if the model is accurate and reliable. It is never clear if simplifications made through analysis are entirely justifiable. There are dangers that new findings are not integrated to the level they might be. The evaluation of policy advice at this level might be seen analogously to the policy itself: matters are too complex to ever know if the best possible has been achieved. Instead, the best which can be done is avoidance of obviously erroneous results and a commitment to continue attempting to improve. This message, the message of Lindblom, is that there is a choice between building models which have a reassuringly straightforward heuristic strategy but little hope of any impact and building models which are complex amalgams of different constraints but are useful in policy contexts. Lindblom's target was abstract modelling but the analysis presented here suggests that the critique can be read equally as critiquing an excessive reliance on the rhetoric of being evidence-based. Agent-based modelling can contribute to a greater focus on validity but this focus does not trivially lead to better support to policy decisions.

BIBLIOGRAPHY

- Agar, M. (2005). Agents in Living Color: Towards Emic Agent-Based Models. *Journal of Artificial Societies and Social Simulation*, 8(1), 18.
- Agar, M. (2004). We have met the other and we're all nonlinear: Ethnography as a nonlinear dynamic system. *Complexity*, 10(2), 16-24.
- Alam, S. J., Meyer, R., Ziervogel, G., & Moss, S. (2007). The impact of HIV/AIDS in the context of socioeconomic stressors: an evidence-driven approach. *Journal of Artificial Societies and Social Simulation*, 10(4), 7.
- Arthur, W. B. (1999). Complexity and the Economy. *Science*, 284(5411), 107-109.
doi:[10.1126/science.284.5411.107](https://doi.org/10.1126/science.284.5411.107)
- Axelrod, R. M. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.
- Bak, P. (1996). *How Nature Works* (1st ed.). New York, NY: Copernicus.
- Barreteau, O., Antona, M., d'Aquino, P., Aubert, S., Boissau, S., Bousquet, F., Daré, W., et al. (2003). Our companion modelling approach. *Journal of Artificial Societies and Social Simulation*, 6(2), 1.
- Barthelemy, O. (2006). *Untangling Scenario Components with Agent Based Modelling: an Example of Social Simulations of Water Demand Forecasts*. (Doctoral dissertation). Manchester Metropolitan University, Manchester, UK. Retrieved February 15, 2010 from <http://cfpm.org/cpmrep163.html>
- Becu, N., Perez, P., Walker, A., Barreteau, O., & Page, C. L. (2003). Agent based simulation of a small catchment water management in northern Thailand

- description of the CATCHSCAPE model. *Ecological Modelling*, 170(2-3), 319-331.
- Boden, R., & Epstein, D. (2006). Managing the research imagination? Globalisation and research in higher education. *Globalisation, Societies and Education*, 4(2), 223 - 236.
- Borenstein, S., Bushnell, J., Kahn, E., & Stoft, S. (1995). Market power in California electricity markets. *Utilities Policy*, 5(3-4), 219-236. doi:[10.1016/0957-1787\(96\)00005-7](https://doi.org/10.1016/0957-1787(96)00005-7)
- Brenner, T. (2001). Simulating the evolution of localised industrial clusters -An identification of the basic mechanisms. *Journal of Artificial Societies and Social Simulation*, 4(3), 4.
- Brenner, T., & Werker, C. (2009). Policy Advice Derived from Simulation Models. *Journal of Artificial Societies and Social Simulation*, 12(4).
- Brown, M., Kauffman, S., Palmrose, Z., & Smolin, L. (2008). Can science help solve the economic crisis? *Edge*, 269. Retrieved February 15, 2010 from http://www.edge.org/3rd_culture/brown08/brown08_index.html
- D'Aquino, P., Le Page, C., Bousquet, F., & Bah, A. (2003). Using self-designed role-playing games and a multi-agent system to empower a local decision-making process for land use management: The SelfCormas experiment in Senegal. *Journal of Artificial Societies and Social Simulation*, 6(3).
- David, N. (2009). Validation and Verification in Social Simulation: Patterns and Clarification of Terminology. In R. Goebel, J. Siekmann, & W. Wahlster (Eds.),

- Epistemological Aspects of Computer Simulation in the Social Sciences* (pp. 117–129). Berlin, Germany: Springer Berlin Heidelberg.
- Dawid, H., & Fagiolo, G. (2008). Agent-based models for economic policy design: Introduction to the special issue. *Journal of Economic Behavior & Organization*, 67(2), 351-354. doi:[10.1016/j.jebo.2007.06.009](https://doi.org/10.1016/j.jebo.2007.06.009)
- Day, C. J., & Bunn, D. W. (2001). Divestiture of Generation Assets in the Electricity Pool of England and Wales: A Computational Approach to Analyzing Market Power. *Journal of Regulatory Economics*, 19(2), 123-141. doi:[10.1023/A:1011141105371](https://doi.org/10.1023/A:1011141105371)
- Downing, T., Moss, S., & Pahl-Wostl, C. (2001). Understanding Climate Policy Using Participatory Agent-Based Social Simulation. In *Multi-Agent-Based Simulation* (pp. 127-140). Berlin, Germany: Springer Berlin Heidelberg.
- Drogoul, A., Vanbergue, D., & Meurisse, T. (2003). Multi-agent Based Simulation: Where Are the Agents? In J. Sichman, F. Bousquet, & P. Davidsson (Eds.), *Multi-Agent-Based Simulation II* (pp. 43-49). Berlin, Germany: Springer Berlin Heidelberg.
- Edmonds, B. (2003). Towards an Ideal Social Simulation Language. In J. Sichman, F. Bousquet, & P. Davidsson (Eds.), *Multi-Agent-Based Simulation II* (pp. 50-53). Berlin, Germany: Springer Berlin Heidelberg.
- Edmonds, B. (2009). *Understanding Observed Complex Systems – the hard complexity problem* (CPM Report No. 09-203). Centre for Policy Modelling Discussion Papers. Manchester, UK. Retrieved February 15, 2010 from <http://cfpm.org/cpmrep203.html>

- Edmonds, B., & Moss, S. (2005). From KISS to KIDS – An ‘Anti-simplistic’ Modelling Approach. In P. Davidsson, B. Logan, & K. Takadama (Eds.), *Multi-Agent and Multi-Agent-Based Simulation* (pp. 130-144). Berlin, Germany: Springer Berlin Heidelberg.
- Entriken, R., & Wan, S. (2005). Agent-Based Simulation of an Automatic Mitigation Procedure. In *Hawaii International Conference on System Sciences* (Vol. 3, p. 92c). Los Alamitos, CA, USA: IEEE Computer Society. doi:[10.1109/HICSS.2005.70](https://doi.org/10.1109/HICSS.2005.70)
- Epstein, J. (2006). *Generative social science: Studies in agent-based computational modeling*. Princeton, NJ: Princeton University Press.
- Etienne, M., Le Page, C., & Cohen, M. (2003). A step-by-step approach to building land management scenarios based on multiple viewpoints on multi-agent system simulations. *Journal of Artificial Societies and Social Simulation*, 6(2), 2.
- Fagiolo, G., Birchenhall, C., & Windrum, P. (2007). Empirical Validation in Agent-based Models: Introduction to the Special Issue. *Computational Economics*, 30(3), 189-194. doi:[10.1007/s10614-007-9109-z](https://doi.org/10.1007/s10614-007-9109-z)
- Florida, R. L. (2002). *The Rise of the Creative Class: And How It's Transforming Work, Leisure, Community and Everyday Life*. New York, NY: Basic Books.
- Frigg, R. (2003). Self-organised criticality - what it is and what it isn't. *Studies In History and Philosophy of Science Part A*, 34(3), 613-632. doi:[10.1016/S0039-3681\(03\)00046-3](https://doi.org/10.1016/S0039-3681(03)00046-3)
- Geller, A., & Moss, S. (2008). Growing QAWM: An Evidence-Driven Declarative Model of Afghan Power Structures. *Advances in Complex Systems*, 11(2), 321-336.

- Geller, A. (2006). *Macht, Ressourcen und Gewalt : zur Komplexität zeitgenössischer Konflikte eine agenten-basierte Modellierung [Power, Resources, and Violence: One the Complexity of Contemporary Conflicts. An Agent-Based Model]*. Zürich, Switzerland: VDF Hochschulverlag an der ETH Zürich.
- Gilbert, N. (2008). *Agent-Based Models*. Los Angeles, CA: Sage Publications, Inc.
- Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the Social Scientist* (2nd ed.). New York, NY: Open University Press.
- Gross, D., & Strand, R. (2000). Can agent-based models assist decisions on large-scale practical problems? A philosophical analysis. *Complexity*, 5(6), 26-33.
doi:[10.1002/1099-0526\(200007/08\)5:6<26::AID-CPLX6>3.0.CO;2-G](https://doi.org/10.1002/1099-0526(200007/08)5:6<26::AID-CPLX6>3.0.CO;2-G)
- Grüne-Yanoff, T. (2009). The explanatory potential of artificial societies. *Synthese*, 169(3), 539-555. doi:[10.1007/s11229-008-9429-0](https://doi.org/10.1007/s11229-008-9429-0)
- Heath, B., Hill, R., & Ciarallo, F. (2009). A Survey of Agent-Based Modeling Practices (January 1998 to July 2008). *Journal of Artificial Societies and Social Simulation*, 12(4), 9.
- Judd, K., & Tesfatsion, L. (Eds.). (2006). *Handbook of computational economics*. Amsterdam, The Netherlands: Elsevier.
- Langford, C., Wood, J., & Ross, T. (2003). The origin and structure of the Calgary wireless cluster. In D. Wolfe (Ed.), *Clusters old and new*. Montreal, Canada: McGill-Queens Press.
- Lindblom, C. E. (1959). The science of "muddling through". *Public administration review*, 19(2), 79-88.

- Lindblom, C. E. (1979). Still muddling, not yet through. *Public administration review*, 39(6), 517-526.
- MacKenzie, D. (2008). *Material Markets: How Economic Agents are Constructed*. Oxford, UK: Oxford University Press.
- Malerba, F., & Orsenigo, L. (2002). Innovation and market structure in the dynamics of the pharmaceutical industry and biotechnology: Towards a history-friendly model. *Industrial and Corporate Change*, 11(4), 667-703. doi:[10.1093/icc/11.4.667](https://doi.org/10.1093/icc/11.4.667)
- Manson, S. M. (2002). Validation and verification of multi-agent models for ecosystem management. In M. Janssen (Ed.), *Complexity and ecosystem management : the theory and practice of multi-agent systems* (pp. 63–74). Cheltenham, UK: Edward Elgar.
- Mayer, I. S., Daalen, C. E. V., & Bots, P. W. (2004). Perspectives on policy analyses: a framework for understanding and design. *International Journal of Technology, Policy and Management*, 4(2), 169. doi:[10.1504/IJTPM.2004.004819](https://doi.org/10.1504/IJTPM.2004.004819)
- Meyer, M., Lorscheid, I., & Troitzsch, K. G. (2009). The Development of Social Simulation as Reflected in the First Ten Years of JASSS: a Citation and Co-Citation Analysis. *Journal of Artificial Societies and Social Simulation*, 12(4), 12.
- Micola, A. R., Banal-Estañol, A., & Bunn, D. W. (2008). Incentives and coordination in vertically related energy markets. *Journal of Economic Behavior & Organization*, 67(2), 381-393. doi:[10.1016/j.jebo.2006.12.007](https://doi.org/10.1016/j.jebo.2006.12.007)
- Midgley, D. F., Marks, R. E., & Cooper, L. G. (1997). Breeding Competitive Strategies. *Management Science*, 43(3), 257-275.

- Morgan, M., & Morrison, M. (Eds.). (1999). *Models as mediators : Perspectives on natural and social sciences*. Cambridge, UK: Cambridge University Press.
- Moss, S. (2008). Alternative Approaches to the Empirical Validation of Agent-Based Models. *Journal of Artificial Societies and Social Simulation*, 11(1), 5.
- Nelson, R. R., & Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.
- Newell, M. W., & Grashina, M. N. (2003). *Project Management Question and Answer Book, The* (1st ed.). New York, NY: AMACOM.
- Pahl-Wostl, C., & Hare, M. (2004). Processes of social learning in integrated resources management. *Journal of Community & Applied Social Psychology*, 14(3), 193-206.
doi:[10.1002/casp.774](https://doi.org/10.1002/casp.774)
- Petrosino, A., Boruch, R. F., Soydan, H., Duggan, L., & Sanchez-Meca, J. (2001). Meeting the Challenges of Evidence-Based Policy: The Campbell Collaboration. *The ANNALS of the American Academy of Political and Social Science*, 578(1), 14-34.
doi:[10.1177/000271620157800102](https://doi.org/10.1177/000271620157800102)
- Pyka, A., & Werker, C. (2009). The Methodology of Simulation Models: Chances and Risks. *Journal of Artificial Societies and Social Simulation*, 12(4), 1.
- Pyka, A., & Fagiolo, G. (2005). *Agent-Based Modelling: A Methodology for Neo-Schumpeterian Economics* (Universitaet Augsburg Institute for Economics Discussion Paper Series No. 272). Augsburg, Germany: Universitaet Augsburg, Institute for Economics. Retrieved February 15, 2010 from
<http://ideas.repec.org/p/aug/augsbe/0272.html>

- Richiardi, M., Leombruni, R., Saam, N. J., & Sonnessa, M. (2006). A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulation*, 9(1), 15.
- Schwoon, M., Alkemade, F., Frenken, K., & Hekkert, M. P. (2008). *A complex systems methodology to transition management* (Innovation Studies Utrecht (ISU) working paper series No. 08-12). Utrecht, The Netherlands: Utrecht University, Department of Innovation Studies. Retrieved February 15, 2010 from <http://ideas.repec.org/p/uis/wpaper/0812.html>
- Simon, H. A. (1996). *The Sciences of the Artificial*. Cambridge, MA: MIT Press.
- Stanislaw, H. (1986). Tests of Computer Simulation Validity: What do they Measure? *Simulation & Gaming*, 17(2), 173-191. doi:[10.1177/0037550086172003](https://doi.org/10.1177/0037550086172003)
- Sterman, J. (1984). Appropriate summary statistics for evaluating the historical fit of systems dynamics models. *Dynamica*, 10(2), 51-66.
- Sun, J., & Tesfatsion, L. (2007). Dynamic Testing of Wholesale Power Market Designs: An Open-Source Agent-Based Framework. *Computational Economics*, 30(3), 291-327. doi:[10.1007/s10614-007-9095-1](https://doi.org/10.1007/s10614-007-9095-1)
- Tesfatsion, L. (2009). Empirical Validation: Agent-Based Computational Economics. Retrieved November 26, 2009, from <http://www.econ.iastate.edu/tesfatsi/empvalid.htm>
- Tesfatsion, L. (2010). Electricity: Agent-Based Computational Economics (Tesfatsion). Retrieved February 15, 2010, from <http://www.econ.iastate.edu/tesfatsi/aelect.htm>

- Troitzsch, K. (2004). Validating Simulation Models. In *Proceedings of the 18th European Simulation Multiconference*. Presented at the ESM 2004. Retrieved from <http://www.scs-europe.net/services/esm2004/pdf/esm-43.pdf>
- Ventosa, M., Baíllo, Á., Ramos, A., & Rivier, M. (2005). Electricity market modeling trends. *Energy Policy*, 33(7), 897-913. doi:[10.1016/j.enpol.2003.10.013](https://doi.org/10.1016/j.enpol.2003.10.013)
- Victor, P. A. (2008). *Managing Without Growth: Slower by Design, Not Disaster*. Cheltenham, UK: Edward Elgar.
- Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2), 8.
- Windrum, P., & García-Goñi, M. (2008). A neo-Schumpeterian model of health services innovation. *Research Policy*, 37(4), 649-672.
- Wolfe, D. (2009). *21st century cities in Canada : the geography of innovation*. Ottawa, Canada: Conference Board of Canada.
- Yücel, G., & van Daalen, E. (2009). An Objective-Based Perspective on Assessment of Model-Supported Policy Processes. *Journal of Artificial Societies and Social Simulation*, 12(4), 3.
- Zeigler, B. P. (1976). *Theory of modelling and simulation*. New York, NY: Wiley.