What Are Models for?

Peter McBurney

Department of Informatics, King's College London, London WC2R 2LS, UK

Abstract. In this paper I discuss some of the purposes and functions of building models, particularly agent-based models, and present a comprehensive list of these purposes and functions. Careful thought and attention is needed when modeling domains containing intelligent entities, which is usually the case for agent modeling. Reflection on the challenges involved in such domains leads me to propose the construction of meta-models, which are models of the relationship between an intended model of the domain and the entities in the domain, when the entities may have access to the intended model or its outputs. Agent-based computing approaches provide disciplined means of specifying, designing, developing and evaluating such meta-models.

1 Introduction

What are models for? Most users and developers of models, in my experience, seem to assume the answer to this question is obvious and thus never raise it. In fact, modeling has many potential purposes, and some of these may conflict with one another. Criticism of modeling efforts or the outputs of those efforts may arise because of mis-perception of the aims and purposes of the modeling activity. Agent-based modeling is usually undertaken for domains having autonomous entities, whether living or organizational, and these entities may also be intelligent. In such cases, a model may exert an influence on the domain being modeled, because the entities in the domain may have, or may seek to have, models of their own. I discuss some of these issues in this paper, starting with the issues of representation and prediction of some real world domain. Many of my reflections apply to any type of modeling, not just agent-based modeling and not even just computer simulation modeling; they also apply to both models developed for business and public policy decisions as well as models for research purposes.

2 Representational and Predictive Functions

2.1 Models and Reality

Most modelers when asked what models are for are likely to answer that they are intended to represent some real phenomenon, some portion of reality. Following Rosen [1] and Hughes [2], we might understand the relationship between the model and reality by means of a sequence of relationships between reality and models:[1]

[1] The three process labels are due to Hughes [2].

M. Cossentino et al. (Eds.): EUMAS 2011, LNAI 7541, pp. 175–188, 2012.

- *Denotation* is the process by which the model (let us call it $M1$) is developed in order to represent some portion of reality (called $R1$).
- *Demonstration* is the working-out or working-through of the model, finding the consequences of the initial states or assumptions either deductively or otherwise, so that we move from some initial model state $M1$ to some consequential state or collection of model outputs, $M2$.
- *Interpretation* is the process of inference from the model consequences $M2$ back to the reality intended to be modeled, perhaps back to an inferred or consequential state of reality, called $R2$.

Demonstration could be undertaken in many different ways: deductive mathematical consequence; physical motions, as in a wind-tunnel or in an orrery (a physical model of planetary relationships and motions); or the flow of communicative interactions between individual agents in a multi-agent model of a society. These various mechanisms are referred to as the *internal dynamics* of the model by Hughes [2]. Morgan [3] argues that such dynamics are typically set going by something external to the model, such as the asking of a *"What if"* question. When run, the model's internal dynamics lead it to some resulting state or to the generation of outputs or consequences, as properties of the model. The modeler then uses these model properties to infer conclusions about the real domain that the model was intended to represent. Such inference from the conclusions or consequences generated by the model back to the reality may be contested. In mainstream economic theory, for example, inference from non-deductive models has usually been regarded as problematic, as discussed in [4,5]. One consequence of this view in mainstream economics has been that research drawing on agent-based models has had great difficulty being published.[2] But since all modeling involves abstraction from reality, even deductive inference is only valid if certain governing assumptions about those aspects of reality abstracted away hold true. In past work [7], we identified 12 inference steps necessary to validly draw conclusions about human carcinogenicity of chemicals based on experimental evidence, of which only one step was statistical inference (i.e., inductive inference). Any inference from a mathematical model using deductive internal dynamics back to reality would require similar contextual inference steps for validity of the interpretation stage.

Robert Rosen understood this fact through the use of a category theoretic model of the modeling process [1]. The three successive processes — denotation, followed by demonstration, followed by interpretation — that connect real state $R1$ to real state $R2$ are only valid if the indirect path they construct between $R1$ and $R2$ via $M1$ and $M2$ mirrors some alternative direct path between $R1$ and $R2$. In other words, modeling the domain will only produce valid inference if the successive stages of denotation, demonstration, and interpretation mirror (in some domain-specific sense) the development of the real phenomenon when it transitions from $R1$ to $R2$.

What exactly does it mean *"to mirror"*? In an influential article in 1953 [8], the economist Milton Friedman argued that models aimed at prediction only need to predict well. They do not need to, and indeed may not be able to, describe the world in state

[2] According to [6], only 8 out of 43,000 recent papers in the top twenty journals in Economics drew upon multi-agent models or computational economics.

$R1$, nor even the actual physical or social processes that take the world from state $R1$ to state $R2$. In other words, a model may have good predictive properties without having good representational or explanatory properties. Newton's theory of gravity, for example, predicted the motion of the planets in solar orbit without providing any (or any reasonable) explanation for gravity. Friedman's famous example also involved Newton: a good player of billiards may be able, based on experience and intuition, predict the likely motion of a billiard ball without any knowledge, let alone any use, of Newton's three laws of motion. Whether or not we accept Friedman's argument depends to a large extent on what we believe a model is for: is it essentially to predict the world or only to represent or explain it?

The game theorist and economist, Ariel Rubinstein, has said that his main purpose in creating models is to better understand and to sharpen his intuitions about the real phenomenon being modeled [9]. How does a model help in understanding? In a series of papers, Mary Morgan and colleagues [3,10,11,12,13,14] have argued that models in economics are idealizations or abstractions of aspects of reality which enable the creation of worlds parallel to the real world. By altering the underlying assumptions or by exploring the internal mechanisms of the model, the modeler is able to explore alternatives to the real world by, for example, asking *What if?* questions, considering counter-factual assumptions, or pursuing alternative development paths. Thus, both the world represented *by* the model and the world *of* the model can be explored. In this view of modeling, Morgan argues, models themselves are not best understood as passive recipients of exploration, but as active participants in the creation of knowledge [3]: by their physical, mathematical, or computational nature, models resist some uses and they facilitate others.

Another issue for representation is that most models typically denote multiple real phenomena, not one particular phenomenon or one state of reality. Newton's equations of motion, although developed with the planets of Earth's solar system in mind, in fact constitute a class of distinct models, and only some (perhaps even only one) denote the actual planets of our Solar system. To model the particular solar system we find ourselves in, various variables, called parameters, need to be instantiated with particular values. This process of instantiation of parameter values so as to match a particular reality is called *model calibration.* The idea that modeling is a process of creating classes of models has been explored in economics by John Sutton [15]. An example of this class-of-models approach is the model of diffusion of agent-software technologies across business networks given in [16]. A related view, due to Trygve Haavelmo [17] and Marcel Boumans [18], is that a model can be seen as an experimental design, and the data used for calibration comprises one experimental outcome of it. Running the model with different input parameters or initial values generates additional experimental outcomes.

Calibration of models and assessment of model predictions assumes that we have some way to measure those aspects of reality that our model purports to represent. This can be problematic, for various reasons. In the case of economic and social domains, the aspects of reality used for calibration or model prediction may be social artefacts: most macro-economic variables (e.g., inflation rates, unemployment levels, etc) do not exist in nature, and are themselves socially-constructed entities. Often their construction is a long and technical process, itself drawing on theories or models of the phenomena in question, and subject to debate and contestation along the way. There are many different operationalizations of the variable called the *supply of money*, for example, and so a modeler may have considerable freedom of choice in calibrating or assessing his or her model against monetary reality. A related issue is the use of *stylized facts* for model calibration and assessment: these are generalizations of reality, also usually developed or mediated through some theory. In economics, the "fact" that *an increase in price leads to a fall in aggregate demand* is an example: in any particular market at any particular time, aggregate demand may or may not fall when price rises, for any number of reasons. So many are the exceptions to this particular "fact" that economists even have a name for the exceptions to it, which are called Giffen Goods. In calibrating or assessing the predictions of a model against stylized facts, a modeler is not using a model to represent reality, but using a model to represent another model of reality.

A further issue with model calibration and assessment arises in theoretical physics and elsewhere: we may have no independent access to the reality intended to be represented by the model other than via the model itself. String theory, for example, seeks to model reality by positing a number of additional spatial dimensions to the three which we humans have experience of. Since we do not have access to these dimensions we cannot independently calibrate string theoretic models against them, nor assess any predictions arising from our models about them.[3] It would seem to me that this situation makes absurd any claims that mathematics is *"unreasonably effective"* in modeling scientific phenomena [20], since how we could tell? Our only way to assess the effectiveness or otherwise of mathematical models of physics is via more mathematics.[4] Likewise, marketing models seeking to predict future consumer purchase intentions can not be calibrated independently of the models themselves, since there are no facts of the matter to assess the model against.

2.2 Intelligent Entities

For the social and policy sciences, the real phenomena represented by a model usually include human individuals or organizations. These entities may be intelligent and may thus act in anticipation of future events.[5] Indeed, models of human societies or human activities may well seek to represent entities — for example, economic agents — who themselves have models of their environment, and who may use these models to guide their actions. What should the modeler assume about the models being used by the entities being modeled? This is a question which most agent-based modelers will face at some point, because agent-based models seek to represent entities in some real domain as separate individuals. One very strange answer to this question is given by a branch of economics, so-called *Rational Expectations Theory*, which assumes, firstly, that all

[3] Perhaps this explains why, at the time of writing, String Theory and M-theory models have yet to generate a single empirically-testable prediction after nearly four decades and thousands of person-years of development [19].

[4] I owe this insight to Stephen Reye.

[5] Focusing mainly on biological and ecological domains, Rosen called such phenomena *anticipatory systems* [1].

economic actors in the model have access to an identical model themselves, and, secondly, that this identical model is the very model being constructed by the modeler. This recursive theory, due originally to economist John Muth [21], has become influential in mainstream economics.[6] To anyone outside mainstream economics, however, these two assumptions are simply bizarre.

Another issue for models of intelligent entities is that the model, or even the fact of modeling, may influence the behavior of the real entities. The most famous example is the Black-Scholes-Merton model of options pricing [23,24]. Prior to the development of this model, trade in options and similar derivative financial products was limited because potential traders were not able to coherently price such products, and there was no agreed theoretical basis for determining such prices. Black and Scholes, and separately Merton, proposed a family of models for options pricing which then led to an explosion in trading of these. To develop their models, the modelers needed to make assumptions about how traders would behave in such marketplaces. Once the models were available, traders, having available no other guide to their behavior, adopted the behaviors assumed by the modelers. Borrowing a term from the philosophy of language for utterances which bring about changes in the world [25], sociologists of economics have called such modeling activities, *performative* [26]: they create the very reality they purport to describe.

Other examples in Economics involve the use of economic game theory to design auction mechanisms, particularly for complex domains such as the combinatorial auctions of PCS radio-frequency spectrum in the USA from 1994 onwards by the US Federal Communications Commission [27,28]. Here one can see the models acting as blueprints for the behaviors of the participants, facilitating some behaviors and precluding others.

Game theory also features in another primary instance of models being performative: the development of western military strategy for nuclear weapons. Game theory models of nuclear warfare provided strategists with the language and conceptual frameworks to identify and explore alternative actions and their likely consequences in this domain [29,30]. One weakness with a game-theoretic view of some interaction is that the actions and strategies suggested by the theory rely on all the participants believing they are playing the same "game", and believing that the other participants also believe this of one another. Philip Mirowski even speculates that in the late 1950s, the USA feared that the leaders of the USSR were not playing the same game as they were, and so American

[6] An interesting question for sociologists of economics is why this theory became influential. One possible explanation is that these assumptions may lead the resulting models to be mathematically tractable more often than do more realistic assumptions. Another possible explanation is that rational expectations theory justifies a particular (conservative) position regarding public policy: any macro-economic policy action will be undermined by pre-emptive, countervailing actions by intelligent economic actors able to second-guess the policy makers. Thus, in this view, it is better for a policy-maker to do nothing than for policies to be subverted pre-emptively. As economist George Stigler once suggested [22], in the market for economic theories, as in any other market, the suppliers of theories produce the theories demanded by those potential consumers of theories who have the money to pay for them; perhaps, then, it should not be surprising that mainstream economic theory has tended to provide support for government policies that favour rich and powerful interests.

political and military leaders embarked on a campaign to persuade the Soviet Union to adopt game theory for nuclear strategy also[7] [30].

2.3 Models and Public Policy

Several features of the relationships between models and domains containing intelligent entities have significant consequences for public policy development and implementation. Two important such features, both related to perfomativity of models, are those of self-fulfilling and self-denying prophecies.[8] If everyone in some system believes that all the others in the system will act in a certain way, each person may then act pre-emptively or in mitigation such that the forecasted behaviour is, or is not, brought about. For example, if there are two alternative routes between two towns, and the majority of drivers hear a radio forecast of heavy traffic on the first route, they may all choose instead to take the second route, thus leading to congestion on the second and lighter traffic on the first; the forecast would therefore be self-denying. Alternatively, if the experience of drivers is such that they tend to dis-believe radio forecasts they hear, then the majority may choose the first route, thus fulfilling the forecast.

As would be expected, these issues become important in matters of public policy, particularly where governments or regulators, by their announcements or their words, communicate forecasts to citizens. In Britain in March 2012, drivers of tankers carrying petrol to service stations indicated that they may stage a strike the following month. A Government minister then announced that car-owners should ensure, in anticipation of any strike, that their petrol tanks were full, an announcement that led to mass and immediate panic buying of petrol. There was thus soon a shortage of petrol caused solely by the widespread and erroneous belief that a petrol shortage was imminent. The fact that panic buying led to garages selling out of petrol of course confirmed public beliefs of an impending shortage.[9]

Because of incidents such as this, most western Governments are very careful about what they announce to their citizens, how they announce it, and when, regarding matters of public health, food safety or national security. Such careful consideration is not uniform across all public policy areas, however. Within economics, there is a widespread belief that the more information is widely known to economic agents, the better will be their economic decision-making, and the better the functioning of the economy as a whole. This belief has led most western central banks to release full information about their decision-making processes for deciding policy on interest rates and on other monetary policy instruments; the Bank of England, for example, openly publishes the minutes of the meetings of its Monetary Policy Committee (albeit with a short delay after each meeting). Clearly, the economic actors concerned with decisions about central bank interest rates are (or, are assumed to be) better informed, and possibly more deliberative

[7] US leaders seem to have done so by stating publicly that game theory was *not* useful for nuclear strategy, as a form of reverse psychology.

[8] These terms were coined by sociologist Robert K. Merton in 1948 [31], who was, interestingly, the father of Robert C. Merton, co-developer of options pricing theory [24].

[9] Funder [32] presents another example, from the last days of the German Democratic Republic in 1988-1989.

decision-makers, than are members of the general public concerned with mad-cow disease or similar health scares. But, as far as I am aware, there is no over-arching theory or model of the relationships between model and domain for public policy domains comprising intelligent, purposeful entities that suggests information should be publicly released by governments in the case of economic domains, and not publicly released in the case of health or security domains. Both these contrasting public policies (release all information readily versus release only limited information carefully) seem to be based merely on implicit, untested assumptions about model-domain relationships and the impacts of additional information on participant behaviours.

3 Mensatic and Epideictic Functions

In addition to representation and prediction, models also serve several other functions in use. Here I discuss two functions which are usually either overlooked or implicit. In their list of three primary roles of business models, for example, Baden-Fuller and Morgan [33] appear to overlook both these roles of models.

The first I term the **mensatic** function, from the Latin word for table, *mensa*. Here the model acts as a vehicle to identify interested stakeholders and to bring them together, around a common metaphorical table. Models for forecasting demand can serve this function internally within companies and with interested outsiders, such as distribution partners [34]. For potential new ventures, particularly those in high-technology industries, business models and plans serve this function with investors and with other potential stakeholders, such as regulators and suppliers. Doganova and Eyquem-Renault present a case study of a potential high-technology start-up in France, describing this mensatic role in detail [35]. In matters of public policy, too, models can act to bring stakeholders together and to help co-ordinate their beliefs and actions. In formulating public health policies for dealing with malaria in developing countries, for instance, epidemiological models can act to co-ordinate the actions of the many stakeholders who need to participate for effective strategy formulation and execution: medical personnel, public health officials, national, regional and local government officials, community and religious leaders, foreign aid donors, international agencies, pharmaceutical companies, and suppliers of other treatment materials. Policy development and planning requires the co-ordination of actions between these various stakeholders in order to design and execute co-ordinated campaigns against the disease [36].

By bringing stakeholders “to the table”, models also serve as the basis for identifying, and potentially deciding, trade-offs in public policy. In complex policy domains such as public health or environmental risk assessment, the potential consequences, costs or benefits of decisions may be experienced differentially by different people or groups within a society, and thus identification of these becomes a major part of public policy formulation [37]. Policy decisions in these domains usually involve complex multi-attribute trade-offs, and, here too, both the making of policy and the forging of a wide public consensus benefit from having different stakeholders discuss and compare decision alternatives [38]. Within western public policy, these deliberative decision processes are probably most finely developed for environmental risk assessment decisions [39,40] and in land-use planning decisions [41].

The mensatic role of models is particularly important for decisions made by multiple people or teams, such as those for major public policy domains. Even for the trading decisions of a private hedge fund, decisions may involve competitors and other outsiders, with the model playing a central role, as Hardie and Mackenize have shown [42]. For economic and marketplace domains, sociologist Michel Callon has called models *market devices*, because they help to engineer, to bring into being, the associated marketplace [43].

Models may also serve an **epideictic** function. Epideictic reasoning involves inference from the form or the style of an argument, rather than from its content only. In an example due to William Rehg [44], suppose you seek advice from two different doctors about treatment for a serious medical condition. One doctor, let us call her Dr X, says that there are three possible courses of treatment. She labels these courses, A, B and C, and then proceeds to walk you methodically through each course – what separate basic procedures are involved, in what order, with what what likely consequences and side effects, and with what costs and durations, what chances of success or failure, and what survival rates. She finishes this methodical exposition by summing up each treatment, with pithy statements such as, *"Course A is the cheapest and most proven. Course B is an experimental treatment, which makes it higher risk, but it may be the most effective. Course C . . . "*, etc.

The other doctor, lets call him Dr Y, in contrast talks in a manner which is apparently lacking all structure. He begins a long, discursive narrative about the many different basic procedures possible, not in any particular order, jumping back and forth between these as he focuses first on the costs of procedures, then switching to their durations, then back again to a discussion of costs, then on to some expected side effects, with tangential discussions about the history of the experimental tests undertaken on one of the procedures, and also about the architect who built the hospital, etc, etc. And he does all this without any indication that some basic procedures are part of larger courses of treatment, or that they are even linked in any way, and speaking without using any patient-friendly labeling or summarizing of the decision-options.

Which doctor would you choose to treat you? If this description was all that you knew, then Doctor X would appear to be the much better organized of the two doctors. Most of us would have more confidence being treated by a doctor who sounds better-organized, who appears to know what he or she was doing, compared to a doctor who sounds disorganized. More importantly, it is also evident that Doctor X knows how to structure what she knows into a coherent whole, into a form which makes her knowledge easier to transmit to others, easier for a patient to understand, and which also facilitates the subsequent decision-making by the patient. We generally have more confidence in the underlying knowledge and expertise of people able to explain their knowledge and expertise well, than in those who cannot, and usually this confidence is justified.

If we reasoned this way, we would be choosing between the two doctors on the basis of their different rhetorical styles: we would be judging the contents of their arguments (in this case, the content is their ability to provide us with effective treatment) on the basis of the styles of their arguments. Such reasoning processes, which use an argument's form to assess its content, are called epideictic, as are arguments which draw attention to their own style.

Since the advent of spreadsheet software applications, business plans for new ventures or new products almost invariably contain a model of the business and of the marketplace in which it will exist. Such business plans and models are often out-of-date very quickly, particularly in turbulent or high-technology markets, or depend on unverifiable conjectures about which there are no facts of the matter (such as future consumer purchase intentions or the reactions of competitors). Investors and other stakeholders, such as distribution partners or suppliers, assessing plans for new business ventures know all this. The function of such business plans and models is not to model or to predict or to control reality accurately, since these goals in any case would usually be impossible. Rather, the function of these models is to force intending new venture managers to engage in structured and rigorous thinking about the domain, and to provide a means by which potential investors in the venture can probe this thinking. By challenging the prior assumptions, the internal dynamics, or the interpretation of the model, potential investors can assess the depth and rigor of the thinking of the management, as well as as assessing managers' flexibility and adaptability in recognizing and responding to changes in the market environment. Investors and other stakeholders thus typically engage in a stress-test of managers' beliefs and plans — contesting the assumptions and reasoning of the business plan; being unreasonable in questions and challenges; prodding and poking and provoking the management team to see how well and how quickly they can respond, in real time, without preparation. In all of this, a decision on the substance of the investment is being made from evidence about the form, of how well the management team responds to such stress testing. This is perfectly rational, given the absence of any other basis on which to make a decision and given our imperfect knowledge of the future. Thus, the business model becomes a vehicle by which potential investors and other stakeholders may assess the capabilities of the management team; the model serves, in other words, an epideictic function.

4 A List of Reasons for Modeling

Several authors have proposed lists of reasons for undertaking modeling, or lists of potential functions of models: Rubinstein [9] lists four reasons for undertaking economic modeling; Bailer-Jones [45] lists five functions of models in science; Epstein [46] lists 17 reasons to build explicit models;[10] and Baden-Fuller and Morgan [33] present three functions of models in business domains. Each of these lists has omissions. Seeking a comprehensive list of reasons for constructing models, I have drawn on these four lists as well as the the reflections above, to create the following list:

1. To understand natural reality: To better understand some real phenomena or existing system. This is perhaps the most commonly perceived purpose of modeling, in the sciences and the social sciences.

2. To predict natural reality: To predict (some properties of) some real phenomena or existing system. As discussed above, a model aiming to predict some domain may be

[10] Epstein's reasons are at multiple levels of granularity, and some of his reasons are the consequences of modeling rather than reasons for doing so, at least for honest modelers, e.g., *"Challenge the robustness of prevailing theory through peturbations"* and *"Expose prevailing wisdom as incompatible with available data"*. He also numbers only 16 of the 17 reasons.

successful without aiding our understanding of the domain at all. For many modeling activities, calibration and prediction are problematic, and so predictive capability may not always be possible as a means of model assessment.

3. To control natural reality: To manage or control (some properties of) some real phenomena or existing system.

4. To understand an existing human model or artefact: To better understand a model of some real phenomena or existing system. Arguably, most of economic theorizing and modeling falls into this category, and Rubinsteins preferred purpose is this type [9].

5. To predict an existing human model or artefact: To predict (some properties of) a model of some real phenomena or existing system.

6. To understand, predict or control a future human model or artefact: To better understand, predict or manage some intended (not-yet-existing) artificial system, so to guide its design and development. Understanding a system that does not yet exist is qualitatively different to understanding an existing domain or system, because the possibility of calibration is absent and because the model may act to define the limits and possibilities of subsequent design actions on the artificial system. The use of speech act theory (a model of natural human language) for the design of artificial machine-to-machine languages, or the use of economic game theory (a mathematical model of a stylized conceptual model of particular micro-economic realities) for the design of online auction sites are examples here. The modeling activity can even be performative, helping to create the reality it may purport to describe, as in the case of the Black-Scholes-Merton model of options pricing discussed above.

7. As a locus for discussion: To provide a locus for discussion between relevant stakeholders in some business or public policy domain, a function I termed, *mensatic*. Most large-scale business planning models have this purpose within companies, particularly when multiple partners are involved. Likewise, models of major public policy issues, such as epidemics, often have this function. In many complex domains, such as those in public health, models provide a means to tame the complexity of the domain. Modeling thus enables stakeholders to jointly explore relevant concepts, data, system dynamics, policy options, and the assessment of potential consequences of policy options, in a structured and shared way.

8. To resolve trade-offs: To provide a means for identification, articulation and potentially resolution of alternative action options, alternative trade-offs, and their consequences in some business or public policy domain; examples include health risk assessment of chemicals or new products by environmental protection agencies, and models of disease epidemics deployed by government health authorities.

9. To structure thinking: To enable rigorous, structured and justified thinking about the assumptions and their relationships to one another in modeling some domain. Business planning models usually serve this purpose. They may be used to inform actions, both to eliminate or mitigate potential negative consequences and to enhance potential positive consequences, as in retroflexive decision making [47].

10. To train people: Models can provide expedited and deliberately-focused experiences of reality, which is why flight simulators are used to train airplane pilots. Market games and marketing models are now commonplace in companies for training of marketing, sales and advertising staff.

11. To assess the modelers: To enable a means of assessment of managerial competencies of the people undertaking the modeling activity. This is the *epideictic* function of modeling, where the model itself is a vehicle to enable interested stakeholders to learn about and assess the assumptions, the reasoning processes, and the future action plans of the people doing the modeling. As mentioned above, business plans and models for new ventures are almost always used in this way by potential investors and business partners to assess the management team of new ventures, and to decide whether or not to participate in the venture.

12. To play: As a means of play, to enable the exercise of human intelligence, ingenuity and creativity, in developing and exploring the properties of models themselves. This purpose is true of that human activity known as doing pure mathematics, and perhaps of most of that academic activity known as doing mathematical economics.

5 Conclusions

The list of reasons for modeling given in this paper shows the diversity of functions of models, particularly when models are created not merely for research, but to support decision-making in business or in public policy. The brief discussion at the end of Section 2.3 about the varying views across different public policy domains on the question of what model information should be available to the entities being modeled, points to the need for the development of meta-models for any model of an important domain. Imagine we seek to model a target domain, X. A meta-model M would include the intended model of the domain, let us call it Model A, together with a representation (another model, B) of the domain X. The key purpose of the meta-model is to better understand (and possibly also to predict and to control) the relationships between Model A and the real intelligent entities inside domain X. Depending on the granularity of our model A, then we may be able to assume that model B is in fact the same model as model A. Likewise, the real entities inside X may be assumed themselves to have access to model A or to model B. As with any model, constructing the meta-model M will allow us to explore *"What if?"* questions, such as alternative policies regarding the release of information arising from model A to the intelligent entities inside domain X. Indeed, we could even explore the consequences of allowing the entities inside X to have access to our meta-model M.

Constructing such a meta-model in any particular domain will not necessarily be straightforward and will require careful thinking and analysis. Because of the recursiveness involved, the thought and analysis required is similar to that used in counterespionage, as described, for example, in [48]. Fortunately, we in the multi-agent systems community have several well-developed techniques for undertaking such meta-modeling: proven methodologies for agent-oriented software engineering, such as *Gaia* [49], and detailed, comprehensive techniques for the careful analysis of dynamic knowledge and belief, such as those in [50]. Arguably, all we currently lack is a good theoretical understanding of joint action, and how it occurs among a group of autonomous agents.

Acknowledgments. This paper arose from an invited talk at the ninth *European Workshop on Multi-Agent Systems (EUMAS 2011)* held in Maastricht, The Netherlands, in November 2011. I thank the organizers for their invitations to speak at EUMAS and to write this paper for the post-proceedings volume. A later version of the talk was given to the 2012 workshop of the *European Working Group on Decision Support Systems (EWG-DSS)*, held in Liverpool, UK, in April 2012. I thank the EWG-DSS organizers for their invitation to speak there, and I am grateful to the audiences of both these talks for their interesting questions and comments. My understanding of the epistemological and pragmatic issues surrounding modeling has benefited greatly from discussions over many years with a large number of people, including Robert Bartels, Martin Chapman, Neill Haine, Tony Hawkins, Benjamin Herd, Peter Lewis, Robert Marks, Marian McEwin, David Midgley, Tim Miller, Lin Padgham, Simon Parsons, the late Ray Paton, Steve Phelps, Stephen Reye, John Roberts, Edward Robinson, Elizabeth Sklar, Elizabeth Sonenberg, Robert Stratton, Nassim Taleb, and the late Ted Wheelwright. The views expressed here, along with any mistakes or omissions, are my own.

References

1. Rosen, R.: Anticipatory Systems: Philosophical, Mathematical and Methodological Foundations. Pergamon Press, Oxford (1985)
2. Hughes, R.I.G.: Models and representation. Philosophy of Science 64(suppl.), S325–S336 (1997)
3. Morgan, M.S.: Models, stories and the economic world. Journal of Economic Methodology 8(3), 361–384 (2001)
4. Marks, R.E.: Validating simulation models: a general framework and four applied examples. Journal of Computational Economics 30(3), 265–290 (2007)
5. Midgley, D.F., Marks, R.E., Kunchamwar, D.: Building and assurance of agent-based models: an example and challenge to the field. Journal of Business Research 60(8), 884–893 (2007)
6. Leombruni, R., Richiardi, M.: Why are economists sceptical about agent-based simulations? Physica A 355, 103–109 (2005)
7. McBurney, P., Parsons, S.: Dialectical argumentation for reasoning about chemical carcinogenicity. Logic Journal of the IGPL 9(2), 191–203 (2001)
8. Friedman, M.: The methodology of positive economics. In: Friedman, M. (ed.) Essays in Positive Economics, pp. 3–43. University of Chicago Press, Chicago (1953)
9. Rubinstein, A.: Modeling Bounded Rationality. MIT Press, Cambridge (1998)
10. Morgan, M.S., Morrison, M.: Models as Mediators. Cambridge University Press, Cambridge (1999)
11. Morgan, M.S.: Symposium on *Marshall's Tendencies*: How models help economists to know. Economics and Philosophy 18, 5–16 (2002)
12. Morgan, M.S.: Imagination and imaging in model building. Philosophy of Science 71, 753–766 (2004)
13. Morgan, M.S.: 1: Modelling as a method of enquiry. Cambridge University Press (2009)
14. Morgan, M.S., Knuuttila, T.: Models and Modelling in Economics (2008)
15. Sutton, J.: Marshall's Tendencies: What can economists know? MIT Press, Cambridge (2000)
16. McKean, J., Shorter, H., Luck, M., McBurney, P., Willmott, S.: Technology diffusion: the case of agent technologies. Autonomous Agents and Multi-Agent Systems 17(3), 372–396 (2008)

17. Haavelmo, T.: The probability approach in econometrics. Econometrica 12(suppl.), 1–115 (1944)
18. Boumans, M.: How Economists Model the World to Numbers. Routledge, London (2005)
19. Smolin, L.: The Trouble With Physics: The Rise of String Theory, the Fall of a Science, and What Comes Next. Houghton Mifflin, New York City (2006)
20. Wigner, E.P.: The unreasonable effectiveness of mathematics in the natural sciences. Communications on Pure and Applied Mathematics 13, 1–14 (1960)
21. Muth, J.F.: Rational expectations and the theory of price movements. Econometrica 29(3), 315–335 (1961)
22. Stigler, G.J.: Do Economists matter? Southern Economic Journal 42(3), 347–354 (1976)
23. Black, F., Scholes, M.: The pricing of options and corporate liabilities. Journal of Political Economy 81(3), 637–654 (1973)
24. Merton, R.C.: Theory of rational option pricing. Bell Journal of Economics and Management Science 4(1), 141–183 (1973)
25. Austin, J.L.: How To Do Things with Words. Oxford University Press, Oxford (1962)
26. MacKenzie, D.A.: An equation and its worlds: *Bricolage*, exemplars, disunity and performativity in financial economics. Social Studies of Science 33, 831–868 (2003)
27. Alexandrova, A.: Connecting economic models to the real world: game theory and the FCC. Philosophy of the Social Sciences 36, 173–192 (2006)
28. Nik-Khah, E.: Designs on the Mechanism: Economics and the FCC Auctions. Ph.D., University of Notre Dame, Notre Dame, Indiana, USA (December 2005)
29. Kahn, H.: On Thermonuclear War, 2nd edn. Princeton University Press, Princeton (1961); First edition published 1960
30. Mirowski, P.: Machine Dreams: Economics Becomes a Cyborg Science. Cambridge University Press, Cambridge (2002)
31. Merton, R.K.: Social Theory and Social Structure. Free Press, New York City (1968)
32. Funder, A.: Stasiland: True Stories from Behind the Berlin Wall. Granta Books, London (2003)
33. Baden-Fuller, C., Morgan, M.S.: Business models as models. Long Range Planning 43, 156–171 (2010)
34. McBurney, P., Parsons, S., Green, J.: Forecasting demand for new telecommunications services: an introduction. Telematics and Informatics 19(3), 225–249 (2002)
35. Doganova, L., Eyquem-Renault, M.: What do business models do? innovation devices in technology entrepreneurship. Research Policy 38, 1559–1570 (2009)
36. The malERA Consultative Group on Modeling: A research agenda for malaria eradication: modeling. PLoS Medicine 8(1), e1000403 (2011)
37. Page, T.: A generic view of toxic chemicals and similar risks. Ecology Law Quarterly 7(2), 207–244 (1978)
38. Fiorino, D.J.: Environmental risk and democratic process: a critical review. Columbia Journal of Environmental Law 14, 501–547 (1989)
39. Graham, J.D., Rhomberg, L.: How risks are identified and assessed. Annals of the American Academy of Political and Social Science 545, 15–24 (1996)
40. Graham, J.D., Wiener, J.B. (eds.): Risk Verus Risk: Trade-offs in Protecting Health and the Environment. Harvard University Press, Cambridge (1995)
41. Forester, J.: The Deliberative Practitioner: Encouraging Participatory Planning Processes. MIT Press, Cambridge (1999)
42. Hardie, I., MacKenzie, D.: Assembling an economic actor: the *agencement* of a hedge fund. The Sociological Review 55(1), 57–80 (2007)
43. Muniesa, F., Millo, Y., Callon, M.: An introduction to market devices. The Sociological Review 55, 1–12 (2007)

44. Rehg, W.: Reason and rhetoric in Habermas's Theory of Argumentation. In: Jost, W., Hyde, M.J. (eds.) Rhetoric and Hermeneutics in Our Time: A Reader, pp. 358–377. Yale University Press, New Haven (1997)
45. Bailer-Jones, D.M.: When scientific models represent. International Studies in the Philosophy of Science 17(1), 59–74 (2003)
46. Epstein, J.M.: Why model? Keynote address to the Second World Congress on Social Simulation. George Mason University, USA (2008)
47. Wohlrapp, H.: A new light on non-deductive argumentation schemes. Argumentation 12, 341–350 (1998)
48. Bagley, T.H.: Spy Wars: Moles, Mysteries and Deadly Games. Yale University Press, New Haven (2007)
49. Wooldridge, M.J., Jennings, N.R., Kinny, D.: The Gaia methodology for agent-oriented analysis and design. Autonomous Agents and Multi-Agent Systems 3(3), 285–312 (2000)
50. Ditmarsch, H., Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Synthese Library, vol. 337. Springer, Berlin (2007)