

Experiments versus models: New phenomena, inference and surprise

Mary S. Morgan

Abstract A comparison of models and experiments supports the argument that although both function as mediators and can be understood to work in an experimental mode, experiments offer greater epistemic power than models as a means to investigate the economic world. This outcome rests on the distinction that whereas experiments are versions of the real world captured within an artificial laboratory environment, models are artificial worlds built to represent the real world. This difference in ontology has epistemic consequences: experiments have greater potential to make strong inferences back to the world, but also have the power to isolate new phenomena. This latter power is manifest in the possibility that whereas working with models may lead to ‘surprise’, experimental results may be unexplainable within existing theory and so ‘confound’ the experimenter.

Keywords: models as mediators, experiments, back-inference, new phenomena

1 THE ROLE OF EXPERIMENTS IN SCIENCE

The conventional belief from outside the experimental community is that the function of experiments in economics is for testing theories. This seems to be understood in the rather strong sense that experiments are conducted to confirm or deny some well specified economic theory. Of course, some experiments are indeed the place where theories can be tested, but to restrict the legitimate role of experiments to just this one function seems unnecessarily limited. An alternative and long-standing tradition in the sciences understands experiment as a creative activity, and an equally long tradition in philosophy of science portrays such activity not only as legitimate, but as one of the main functions of experiment. For example, Ian Hacking (1983), in a well-regarded contemporary account, suggests that it is scientific curiosity, first cousin to Veblen’s ‘idle’ curiosity in technology, which drives the experimentalist. Hacking explores a variety of ways in which careful observation and experiment leads theory, and concludes ‘... although some theory precedes some experiment, some experiment and some observation precedes theory, and may for long have a life of its own.’ (1983: 160).

Hacking's account of experiment is based around the notion of isolating phenomena, where phenomena are understood as noteworthy, discernible events that 'occur regularly under definite circumstances' (1983: 221). Experimentally isolated regularities are *noteworthy* precisely because they don't fit into current theoretical accounts (they neither confirm nor deny the theory), but rather offer sufficiently recognisable and ordered outcomes, regular patterns, that beg for description and consequent interpretation.¹ With regularity comes the possibility to 'solve' the phenomena – to provide descriptive laws, although not necessarily to provide a causal account. One obvious example of the use of economic experimentation to 'create, produce, refine and stabilize phenomena'² (1983: 229–30) is the famous case of preference reversals. Here, the experimental work led to new regularities observed in the characteristics of economic behaviour, new classifications and thence to new concepts, and theorising to explain these regularities.

Following this tradition in which experiments are understood as an explorative creative activity brings an immediate contrast, and possible comparison, with the role of mathematical models in economics. As Daniel Hausman (1992) has argued, one of the main role of models in economic theorising is in creative work in the conceptual activity of classifying and describing phenomena. I heartily endorse this view, but how does it do so? In previous analyses (see Morgan 2002, 2003), I have suggested that the method of exploring the mathematically expressed 'world in the model' is a kind of experimental activity that enables the economist to develop theories and explore the limits and implications of the theories manifest in the model.³

Experiments and models have much in common in the way they are used in economics. They share traits which enable both to operate as epistemic mediators: tools of investigation to help find out about either the world or theories. These claims of similarity are not uncontentious, so I begin with a comparison, starting out from the models side.

2 EXPERIMENTS AND MODELS: CONSTRUCTION

The usual question asked about models in the philosophy of science has, until recently, been: What are models? In contrast to the literature on the role of experiments, little attention had been paid to the question 'What roles do models play in science?' In *Models as Mediators* (Morgan and Morrison 1999, ch. 2) we sought answers to how models function by starting from the observation that models act as autonomous instruments of investigation. This starting point was in itself something of a break with the standard accounts that portrayed models either as versions of theories or as versions of data and so lacking this autonomous status.

As with all technologies, the power of models and experiments as instruments of investigation is only evident in the context of their use, but

their use depends in the first instance upon their construction. In *Models as Mediators*, we argued that models are constructed from elements of theory, elements of data, analogical elements, and so forth. Experiments seem to have the same kind of construction. Theories do not provide directly for experiments any more than observations do. Theories and observations, background and foreground knowledge about the world, are resources that go into the design of an experiment. Sometimes there is more of theory, sometimes less.

Marcel Boumans (1999) makes this construction account explicit in his analogy of model construction being like cake-baking according to a new, or to an already developed, recipe. Like Boumans' model recipes, sometimes experiments are newly designed, but much more often they are variations on a well-understood, even classic, design. Rainier Lange (2003) adopts the same metaphor for experiments, seeing 'Experimental instructions as a special kind of recipe' to 'produce, control and register the phenomena that their [scientists'] publications refer to' (2003: 120, 119) and he too suggests that these recipes are 'conspicuously stable under theory change' (2003: 123). There is rarely a completely newly designed experiment or a newly constructed model, although a new recipe in either genre may supply the potential to redirect the field.

In economic experiments, as in all experiments, much is taken from standard rules about experimental design in the field. This has its equivalent in modelling, in, for example, habitual assumptions made for mathematical tractability or in the mathematical formalism chosen to mould the elements in the model together (see Boumans 1999 again), although these are surely less rule-like than the experimental equivalents.

It is particularly these independent elements – analogies, rules, mathematical forms etc, that is, elements which come into construction and are independent of either theory or the world – which may allow experiments, like models, to work autonomously. It is their degree of autonomy which enables models, like experiments, to gain 'a life of their own'. That is, like experiments, certain models begin to take on a life independent of their original creation. Samuelson's overlapping generations model is a good example (see Hausman 1992). Hicks IS/LM model is another such, running for more than 50 years, but soon coming free from its original demonstration purpose (to show the difference between Keynes' theory and classical economics) and from its subsequent Keynesian theoretical orientation. Similarly, Chamberlin's (1948) and Smith's (1962) classroom market experiments created a classic recipe.⁴ Their experimental design has been adopted, adapted, gone through many variations, and been extended for new purposes over the years, so that it has come to have an autonomous life, cut free from the initial experimenters and their experiments.

However, partial independence in construction is not the only requirement for creative work and for learning in the 'models as mediators'

comparison. The second hinges on the characteristic trait that models are constructed to incorporate some form of representation: of theory, of the world, or possibly of both at once, so that in using the model, something may be learnt about the representation which provides for inference about the theoretical or world elements. In the comparison here then I need to look for the equivalent to this ‘representing’ quality in the case of experiments. Here ‘representing’ needs to be replaced by ‘replicating’, in the sense of ‘reproducing’ or ‘capturing’ some aspect of the world in the same kinds of materials and forms as those of the real economic world. In other words, the construction of an experiment recreates a part of the real world inside the artificial environment of the laboratory.⁵ In contrast, the construction of a model creates an ‘artificial’ world – artificial because made out of mathematics, diagrams, or alternative physical domain materials (i.e., not those of the economics domain). As it will turn out, this matter of their materials – the artificial world vs. the real one of people and economic decisions – is an important element in the comparison of the status of a model compared to that of an experiment.

3 FUNCTIONING SIMILARITIES: EXPERIMENTS AS MEDIATORS

The partial independence of construction suggested in the models as mediators approach is the basis of the autonomous mediating role we proposed for models. If experiments are constructed like models, and function autonomously rather like them, they might also act as flexible mediators between theories and the world in the way that models do. (Indeed, Guala 1998 has already suggested such a role for experiments.⁶) The second part of this comparison, then, concerns the way models and experiments function.

Economists’ manipulations of models can be understood as a kind of experiment, although there are also some obvious differences between the modes in which model experiments and laboratory experiments work, as shown in Table 1. As the comparisons in the Table show (rows 3 and 4), models institute experimental control by a combination of model design and *ceteris paribus* assumptions,⁷ experiments do so by a combination of design and real control; model experiments demonstrate their results by methods of mathematical manipulation, experiments do so by laboratory manipulation. This Table extends my earlier analysis to clarify how these different kinds of experiments work in economics (see Morgan 2002 and 2003 for a full analysis of these cell entries, and further comparisons with simulation).

Two further important comparisons come out in the next two rows showing relationships to subject matters and issues of inferences. (The elements of danger and surprise, will be discussed later.) The presumption in the conventional belief about the theory-testing function of experiments in

Table 1 Experiments and models compared

	Mathematical model experiment	Ideal laboratory experiment
Construction and design	Create an artificial world in the model	Re-create part of the real world in an artificial environment
Relationship to subject matters	<i>Represent</i> theory or world	<i>Replicate</i> theory or world
Experimental control	By assumptions of model design and <i>ceteris paribus</i> conditions	By experimental design and physical controls
Demonstration method	Deductive/ mathematical in model	Experimental in laboratory
Inference to world	Different materials. Problems of 'realism'	Same materials. Problems of 'validity'
Inference to theory	Same materials	Different materials
Design danger	'Built-in justification'*	Over-tamed behaviour
Results potential	Surprise	Surprise and confoundment

* This derives from Boumans' 1999 discussion.

economics is that the theory under test is embedded in some way in the experimental design and that running the experiment under the experimental controls demonstrates some outcomes which offer a test of the theory's claims against the behaviour manifest in the experimental world. From this it may be possible to learn something about the theory's applicability to that experimental world.

This account of theory testing fits neatly into an 'experiments as mediators' framework. The power of the experiment to test the theory in this way depends upon the experimental set up satisfactorily capturing a part of the world and accurately using the relevant theory too, so that inference about the theory can be carried through to say whether it applies to the experimentally defined world or not.⁸ Of course, there are further very important questions here about whether a theory supported by the experimental results transfers beyond the experimental set up into the wider world, a problem recognised by economic experimentalists as 'parallelism' or 'external validity' (see Guala forthcoming). Of course, the validity issue is also connected with another point shown in the Table analysis, namely the method of experimental controls used. Clearly any comparison with the model experiment is still very much to the real experiment's advantage here, for the difficulty of domain assumptions in laboratory experiment, where some form of real economic behaviour is under experiment, is nothing

compared to the equivalent well-known problem of the realism of the assumptions in model experiments which involve an artificial world made of different materials. But these points take us away from the creative, exploratory role of experiment, and the associated forms of inferences to behaviour in the world.

4 ONTOLOGICAL CLAIMS AND EPISTEMOLOGICAL POWER

The Baconian tradition of using experiment was epitomised in Hacking's terms as a way of 'creating, producing, refining and stabilizing phenomena' within a technology. Recent philosophers of science have paid more attention to the role of such technologies, particularly Rom Harré, who distinguishes between two kinds of laboratory apparatus. One kind, particularly relevant here, involves an instrument/world mix which serves 'as a working model of some part of the world' (2003: 26). Note that 'model' does not mean a mathematical model. He is writing here about material models, some of which function as 'domesticated versions of natural systems' (2003: 26) (e.g., a fruit fly colony in the lab or an Atwood machine for dropping weights). My notion of experiments as capturing or reproducing the world is well matched in Harré's label of domesticated versions of natural systems, for the important point for both of us is that these are made of the same stuff as the real world. His domesticated model systems match my laboratory worlds, they are of the same materials as the world, yet 'simpler, more regular, and more manipulable' (2003: 27) than those in the wild.

As Harré notes, phenomena are created in these domesticated model worlds that do not happen in the wild (e.g., mutant flies in the lab that would not occur or would not live in the wild). But these kinds of domesticated worlds, and the phenomena that they create, are *qualitatively* different from another kind where the world and the lab equipment jointly 'bring into existence phenomena that do not exist as such in the wild' (2003: 28). This second kind of apparatus he calls 'Bohrian complexes', and the phenomena they generate are 'properties of a complex unity, the apparatus/world entity'. The obvious scientific example for the Bohrian apparatus/world complex type are the particle phenomena produced at CERN, but a homely (technological, rather than scientific) example is a loaf of bread: 'A loaf is brought into existence from wheat and other ingredients by the use of material structures that do not exist in the wild, such as flour mills and ovens. Loaves do not appear spontaneously in nature' (2003: 30). However, the line dividing these two kinds of apparatus is not well defined.⁹ Some might seem obvious just because (some think) there is no such system in existence in the wild, e.g., CERN or cloning technologies. Others seem less clear-cut: in-vitro fertilisation could be understood as a domesticated model

world, but certain techniques used with it, e.g., kinds of gene therapy, might turn it into a Bohrian complex.

What is the relationship between the phenomena created in the domesticated world experiment and the real world? Here Harré has the useful term of ‘back inference’ meaning inference from the experiment back to nature or the wild world:

Domestication permits strong back inference to the wild, since the same kind of material systems and phenomena occur in the wild and in domestication. An apparatus of this sort is a piece of nature in the laboratory. Of course, the richness of back inference will depend on how relations of similarity and difference are weighted by the interests of the researcher in performing the experimental manipulations. There is no ontological disparity between apparatus and the natural setup. The choice of apparatus and procedure guarantees this identity, since the apparatus is a version of the naturally occurring phenomenon and the material setup in which it occurs. (2003: 27–8).

As my Table cell entries suggest, the fact that the same materials are in the experiment and the world makes inferences to the world possible if not easy. As Harré lays it out, it is clear why it is a strong form of inference: the shared ontology has epistemological implications. We are more justified in claiming to learn something about the world from the experiment because the world and experiment share the same stuff. In contrast, inference from the model experiment is much more difficult as the materials are not the same – there is no shared ontology, and so the epistemological power is weaker.

What follows from this for economics’ experiments? While it is certainly tempting to understand economic experiments as using apparatus of the domesticated model world kind, much hangs on their labelling as such. If economic experiments are always on domesticated worlds, then back inference possibilities are stronger and cleaner-cut than for Bohrian complexes, where, as Harré points out, we would have a much more difficult task.

5 SURPRISE VERSUS CONFOUNDMENT

I have been considering the following questions: Are experiments and models the same kind of investigative instrument? The answer seems from the above to be ‘Yes’ in that they share similar structures in construction (as Mäki, this volume, also argues) and they have the same autonomous status which enables them to function in rather similar ways as instruments of investigation. But the answer (contra Mäki), is also ‘No!’ In spite of their similarities models are not experiments, for they do not have the same power as epistemic mediators.

The critical difference lies in the materials: economics is about people's behaviour in the real world, and economic experiments are experiments directly on that real behaviour, however constrained and squashed to fit into the laboratory world that behaviour might be. The model experiments run by economists are experiments with other materials, mainly pieces of mathematics. This materials difference matters greatly for two reasons. Harré used a comparison between different kinds of laboratory experiment to explain the first reason: the ontology matters because it affects the power of inference. The second reason relates to the role of experiments in creating new phenomena and stems from my comparison with mathematical models. This difference in potential is recognised in my notion that where mathematical modelling results may *surprise*, experimental results may not only surprise but also *confound* (see Morgan 2003).

Within the community of experimentalists, during seminars and conference presentations, the liveliest of discussions come in speculating, theorising, and in other ways trying to explain the surprising behaviour patterns made manifest in experiments. If, after many experimental replications with many subjects and with slight variations in the experimental design, certain experimental behaviours continue to be manifest (e.g., co-operation in prisoner dilemma games, preference reversals, and so forth), then the interpretation of the behaviour changes. It is no longer regarded as an experimental artefact and becomes a genuine behaviour pattern, and instead of trying new experimental designs to get rid of the artefact the focus becomes one of explaining the pattern. Such reactions are more evident and more speedy inside the experimentalists' community than in the public sphere of journal articles, but eventually the transition is made and the behaviour pattern comes to be accepted as a new phenomenon.

Such new behaviour patterns, ones that surprise and at first confound the profession, are only possible if experimental subjects are given the freedom to behave other than expected by the experimenter-economist. This is an important consideration in the design of experiments: experiments need to be set up with a certain degree of freedom on the part of participants so that their behaviour in the experiment is not totally determined by the theory involved, nor by the rules of the experiment. If the experimental behaviour is totally predetermined, there is no potential for unexpected patterns to emerge. Where there is no potential to exhibit alternative and unexpected behaviour, there can be no true potential to confirm theory or to refute it. There must be potential to confound the experimenter with noteworthy results which are both surprising and unexplainable within the given realm of theory.

This potential for laboratory experiments to surprise and confound contrasts with the potential for mathematical model experiments only to surprise. In mathematical model construction, the economist knows the resources that went into the model. Using the model may reveal some

surprising, and perhaps unexpected, aspects of the model behaviour. Indeed, the point of using the model is to reveal its implications, test its limits and so forth. But in principle, the constraints on the model's behaviour are set, however opaque they may be, by the economist who built the model so that however unexpected the model outcomes, they can be traced back to, and re-explained in terms of, the model.

It is this potential for independent action by the experimental participants from which new phenomena emerge that supports the greater epistemic power of laboratory experiments over mathematical model experiments. In laboratory experiments, much is determined and constrained by the design set-up and rules of the experiments. This is clearly as it should be – there must be experimental controls, just as a mathematical model is limited by assumptions (about for example, starting values, parameter values and *ceteris paribus* assumptions). Laboratory experiments are not like field experiments – experiments that follow economic behaviour 'in the wild'.¹⁰ Laboratory experiments, by their design, investigate tamed economic behaviour – tamed, or domesticated (to use Harré's term), by the experimental world in the laboratory, that is by the experimental design and experimental rules that constrain behaviour. However, if the behaviour of those taking part in the laboratory experiment is entirely constrained, then the results will be determined absolutely by the experimental design and rules. Particularly when models are used within an experiment to provide measures of control, they may tighten the theoretical constraint beyond usefulness. They raise the danger of over-taming the participants in a particular way so that participants are no longer domesticated, but agents whose behaviour is directed by models of the world, models dictated by the economist.¹¹ In such cases, as in mathematical model experiments, the experimenter may be surprised, but not actually confounded. An important question for any experimental design therefore is: Where is the potential for independent action in the experiment?

Consider, as an exemplar of these points, the classic early classroom experiments on how markets work by Edward Chamberlin (1948), in which participants acting as buyers or sellers were given cards with reservation prices that were chosen by the experimenter as numbers taken from a supply and demand model with pre-specified parameters. The numbers generated by the model, and the 'rules' of the market, provided some constraint on the subject's behaviour but also allowed them freedom to trade at some range of prices. The question for the experiment was whether the subjects, given the rules of bilateral negotiation, and, to them, the unknown market relations model, would exchange at the 'equilibrium' prices and quantities predicted by the normal economic Marshallian solution, or determined by 'the auctioneer' in the Walrasian solution, or whether they would do something else. In Chamberlin's experiments, he found that the average price was lower

than predicted and sales higher than predicted, and explored the reasons for this unexpected behaviour.

To Chamberlin, the recalcitrance of his results – the phenomena of a pattern of behaviour – seemed not just surprising, but sufficiently unexpected and startling enough to make him think seriously about a fundamental assumption in microeconomics. On the basis of his results, he came to doubt that there was even a tendency towards equilibrium. The danger here is manifest and Chamberlin avoided it. Behaviour over-tamed by the design or rules to ensure that participants behaved according to the model would have produced equilibrium convergence: it could not have tested such a theory about markets, and there would have been none of the surprise or confoundment that Chamberlin experienced.

6 CONCLUSION

Economists have in the past seemed to be unduly suspicious of experiments and perhaps over willing to accept that modelling is the best way to do economics. The intuition of other sciences and philosophy is that experiment is a more reliable guide to scientific knowledge. In this comparison of experiments and models, I have argued that on grounds of inference, experiment remains the preferable mode of enquiry because ontological equivalence provides epistemological power. But the importance of experiment in the creation or isolation of new phenomena might well be more critical. Hacking's recipe for the social sciences is that scientific advance requires speculation, calculation and experiment in collaboration with real theoretical entities, 'entities which are part of the deliberate creation of stable new phenomena' (1983: 249). This recipe seems much more likely to be realised from the experimental community than the modelling community.

Mary S. Morgan
London School of Economics and University of Amsterdam
m.morgan@lse.ac.uk

ACKNOWLEDGEMENTS

This paper draws on two previously published papers (based on research funded by the British Academy, whom I thank) and materials from a forthcoming book. A preliminary version was given at the 'Workshop in Experimental Economics' in Nottingham, UK, September 2003 and a more advanced version at the Amsterdam INEM meeting in August 2004. I thank Bob Sugden, Dan Hausman, and the many LSE, Amsterdam and other

colleagues, particularly Marcel Boumans with whom I first worked through this comparison and Francesco Guala who has been entirely generous with his comments.

NOTES

- 1 Hacking's notion that one of the functions of experiments is as an empirical means to isolate phenomena also appears in Mäki's (1992) account of economic modelling as a process of theoretical isolation, but without the former's attention to new phenomena. Sugden's account (in this volume) of 'exhibits' appears rather closer to Hacking's notions.
- 2 Note: this is not a radical social constructivist account: nature plays a role in the production of phenomena, but it is only with a particular technology – the repeated experiment – that we can extract the regularity and present it as 'pure, isolated, phenomena' (Hacking 1983: 226).
- 3 Although it is not their only role, as Gibbard and Varian pointed out as long ago as 1978, and Robert Sugden (2002) and I have separately discussed (Morgan 2001), mathematical models are also applied in casual manner to explain phenomena in the world.
- 4 For example, see Kagel and Roth 1995: 49–55. As Roth remarked, in his survey introduction, two chapters in this *Handbook of Experimental Economics* testify to the 'continuing importance of Chamberlin's 1948 experimental design' (49).
- 5 In this volume Mäki argues for a really rather different view, namely that models and experiments share a similar representing relation to the world: one of 'resemblance', while Bardsley discusses the argument that the attempt to capture the world in the laboratory alters it to the point of artificiality.
- 6 See also his contribution to this volume. Space constraints have limited my references to Guala's many papers on different aspects of the philosophy of economic experimentation; for references, see his forthcoming book (2005).
- 7 'Model design' includes simplification and idealisation assumptions which provide some kinds of control, for example, that certain relations are independent, while a variety of *ceteris paribus* assumptions mark out that certain things are zero, others negligible, others constant, etc; see Boumans and Morgan 2001.
- 8 Such a theory-testing account of experiment seems to portray experiment as lying unambiguously in the middle, between theory and the world rather as in the simulacrum account of models of Cartwright (1983). It may also be consistent with Hausman's (1992) methodological account of models in economics, that we test a theory by claiming that a particular model that offers an interpretation of the theory is, or is not, true of the world.
- 9 As indeed Henk de Regt argued at the Amsterdam conference on experiments, where Harré's paper on domestications and mine on the model/experiment comparisons were first presented (see Radder 2003).
- 10 Only field experiments can investigate behaviour in the wild where controls function not to control the behaviour via *ceteris paribus* assumptions (of mathematical model experiments) nor the *ceteris paribus* controls (of the laboratory world limitations) but by careful physical designs (such as Latin Squares), or statistical randomisation, to wash out the effects of the lack of *ceteris paribus* controls.

- 11 This over-tamed behaviour has its counterpart in the world of mathematical model experiments by the problem of ‘built-in justification’ as defined by Boumans (1999), where sufficient elements of empirical behaviour are built into the model at its construction stage resulting in the justification of the model when it is used. So, model builders too can limit their surprise factor, and thus potential for learning.

REFERENCES

- Boumans, M. (1999) ‘Built-in justification’, in M. Morgan and M. Morrison (eds) *Models as Mediators: Perspectives on Natural and Social Science*, Cambridge: Cambridge University Press.
- Boumans, M. and Morgan, M.S. (2001) ‘*Ceteris paribus* conditions: materiality and the application of economic theories’, *Journal of Economic Methodology* 8(1): 11–26.
- Cartwright, N. (1983) *How the Laws of Physics Lie*, Oxford: Clarendon Press.
- Chamberlin, E.H. (1948) ‘An experimental imperfect market’, *Journal of Political Economy* 56(2): 95–108.
- Gibbard, A. and Varian, H. (1978) ‘Economic models’, *Journal of Philosophy* 75: 664–677.
- Guala, F. (1998) ‘Experiments as mediators in the non-laboratory sciences’, *Philosophica* 62: 901–18.
- Guala, F. (2005 forthcoming) *The Methodology of Experimental Economics*, Cambridge: Cambridge University Press.
- Hacking, I. (1983) *Representing and Intervening*, Cambridge: Cambridge University Press.
- Harré, R. (2003) ‘The materiality of instruments in a metaphysics for experiments’, in H. Radder (ed.) *The Philosophy of Scientific Experimentation*, Pittsburg: Pittsburg University Press, pp. 19–38.
- Hausman, D.M. (1992) *The Inexact and Separate Science of Economics*, Cambridge: Cambridge University Press.
- Kagel, J.H. and Roth, A.E. (1995) *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
- Lange, R. (2003) ‘Technology as basis and object of experimental practices’, in H. Radder (ed.) *The Philosophy of Scientific Experimentation*, Pittsburg: Pittsburg University Press, pp. 119–37.
- Mäki, U. (1992) ‘On the method of isolation in economics’, in C. Dilworth (ed.) *Idealization IV: Intelligibility in Science*, Amsterdam: Rodopi, pp. 317–51.
- Mäki, U. (2002) *Fact and Fiction in Economics*, Cambridge: Cambridge University Press.
- Morgan, M.S. (2001) ‘Models, stories and the economic world’, *Journal of Economic Methodology* 8(3): 361–84 (also in Mäki 2002; 178–201).
- Morgan, M.S. (2002) ‘Model experiments and models in experiments’, in L. Magnani and N.J. Nersessian (eds) *Model-Based Reasoning: Science, Technology, Values*, New York: Kluwer Academic/Plenum Publishers, pp. 41–58.
- Morgan, M.S. (2003) ‘Experiments without material intervention: model experiments, virtual experiments, and virtually experiments’, in H. Radder (ed.) *The Philosophy of Scientific Experimentation*, Pittsburg: Pittsburg University Press, pp. 216–35.
- Morgan, M.S. and Morrison, M. (1999) *Models as Mediators: Perspectives on Natural and Social Science*, Cambridge: Cambridge University Press.

- Radder, H. (2003) *The Philosophy of Scientific Experimentation*, Pittsburgh: Pittsburgh University Press.
- Roth, A.E. (1995) 'Introduction to experimental economics', in J.H. Kagel and A.E. Roth (eds) *The Handbook of Experimental Economics*, Princeton: Princeton University Press, pp. 1–109.
- Smith, V.L. (1962) 'An experimental study of competitive market behaviour', *Journal of Political Economy* 60(2): 111–37.
- Sugden, R. (2002) 'Credible worlds: the status of theoretical models', in U. Mäki (ed.) *Fact and Fiction in Economics*, Cambridge: Cambridge University Press, pp. 107–136.