

# Models, Simulations, and Experiments

Francesco Guala

*Centre for Philosophy of the Social Sciences, University of Exeter, UK, f.guala@ex.ac.uk*

**Abstract:** I discuss the difference between models, simulations, and experiments from an epistemological and an ontological perspective. I first distinguish between ‘static’ models (like a map) and ‘dynamic’ models endowed with the capacity to generate processes. Only the latter can be used to simulate. I then criticise the view according to which the difference between models/simulations and experiments is fundamentally epistemic in character. Following Herbert Simon, I argue that simulations merely require the existence of an abstract correspondence between the simulating and the simulated system. In experiments, in contrast, the causal relations governing the experimental and the target systems are grounded in the same material. Simulations can produce new knowledge just as experiments do. But the prior knowledge needed to run a good simulation is not the same as that needed to run a good experiment. I conclude by discussing ‘hybrid’ cases of ‘experimental simulations’ or ‘simulating experiments’.

## 1. INTRODUCTION

Empiricist philosophies of science tend to draw a sharp distinction between descriptive or representational devices (scientific theories) and what is described or represented (the natural or social world). Models and simulations are customarily placed among the representational tools, whereas experiments are considered as parts of the natural or social world that have been carefully designed in order to answer some specific question. This framework is certainly useful and captures many aspects of scientific practice. There are, however, bits of science that do not fit neatly, and for which a different scheme of classification may be more appropriate. In this paper I shall try to show that it is sometimes useful to think of models, experiments and simulations as tokens of the same kind, somehow located between our statements about the world (call them scientific laws, principles, theories, axioms), and the world itself (see also Guala, 1998). Borrowing from Margaret Morrison and Mary Morgan (1999), we may say that such entities ‘mediate’ between theory and reality.

First, let us notice that everyday scientific talk does often treat experiments, models and simulations as tokens of the same kind. In one of the earliest papers in the field of experimental economics, for example, the

term ‘simulation’ appears three times only in the first page (see Smith, 1991, p.8), alongside other expressions such as ‘experiment’ and ‘experimental game. Or take medicine. Experimental physiologists make extensive use of animals in their investigations, for well known (although controversial) ethical reasons. Most often, these activities fall under the label of ‘animal experimentation’. But it is not uncommon to hear or read the expression ‘animal models’, especially when the experimenters fear that the findings will not be easily transferable from animal subjects to human beings.

Why do scientists slip from ‘experiment’ talk, to ‘model’ and to ‘simulation’ talk? A plausible answer is that the difference is purely epistemic in character: ‘experiment’ and ‘theory’ being the pillars upon which all proper science should stand, scientists signal their epistemic doubts using a special terminology. An incomplete or less than certain theory becomes a ‘model’; a dubious experiment becomes a ‘simulation’, and so on. However, perhaps this is not the whole story, as I shall try to show in the rest of this paper.

## 2. MODELS AND SIMULATIONS

Models have been at the forefront of research in the philosophy of science during the last couple of decades. Indeed, the latest orthodox ‘theories of scientific theories’, the so-called ‘Semantic View’ of theories, *identifies* theories with sets of models. The Semantic View is more a family of doctrines than a single, unified philosophical theory, but all its versions share a distaste for the older ‘syntactic’ approach, according to which theories are basically sets of statements or laws. In the semantic approach the fundamental component of a theory, the model, is in contrast a *structure* – a set of objects with properties and relations among them and/or their parts – that satisfies the linguistic components of the theory. The latter are secondary, in the sense that they can be formulated in various equivalent ways, as long as they are satisfied by the set of models that make the theory. The axioms, laws, etc., may change depending on the language and system of axioms scientists choose, but the models won’t. The models must be put at work by means of a ‘theoretical hypothesis’, stating that they stand in a certain relation (of similarity, isomorphism, analogy, etc., depending on which version of the Semantic View one subscribes to) with real-world entities or systems. Since the Semantic View is presently the received explication of the concept and role of scientific models, I shall take it as my point of departure here. The next question to be tackled is: what is a simulation?

Mario Bunge (1969) defines simulation as a relation between two

entities,  $x$  and  $y$ , where  $x$  simulates  $y$  if (1) there exists a correspondence relation between the parts or the properties of  $x$  and  $y$ ; and (2) the analogy is valuable to  $x$ , or to another entity ( $z$ ) that controls  $x$  (p.20). The first striking feature of this definition is its anthropocentrism. It makes no sense to say that a natural geyser ‘simulates’ a volcano, as no one controls the simulating process and the process itself is not useful to any one in particular. I shall assume for the sake of this paper that the second part of Bunge’s definition captures some important intuitions and connotations of the term simulation. But the first part is unsatisfactory, because it leads to include things that we would not intuitively call ‘simulations’ at all. Consider a map: if it has been drawn adequately, there must exist some correspondence relation between its parts and parts of the territory it is aimed at representing.<sup>1</sup> Since the map is also somehow ‘controlled’ by someone (its user or its creator), and is certainly valuable to her, it does fulfil all of Bunge’s criteria. Yet, it would seem odd to say that a map ‘simulates’ the territory.

Now consider a map together with a set of flags and miniaturised soldiers and tanks, of the sort you find in military head-quarters or in games such as ‘Risk’. If the toy-flags, mini-soldiers and mini-tanks are moved on the map according to the appropriate rules, we can properly claim that a battle or a military campaign is being simulated. Why? Whereas the map alone is somehow ‘inert’, the same map, plus the miniatures, plus the players or officials moving the miniatures according to the rules, make a ‘dynamic’ system. I shall here follow Stephan Hartmann (1996) and distinguish *static* from *dynamic* models. A static model can only represent a system at rest. A dynamic model can also represent the time-evolution of the system (p.82).<sup>2</sup> A dynamic model, then, can be in different states at different times, and usually each state will correspond to a specific combination of values taken by the variables in the model. Such a model then will be able to be in as many different states as all logically or physically possible permutations of the values its variables can take. Only ‘dynamic’ systems of this sort can properly speaking simulate. “*A simulation imitates one process by another process*” (Hartmann, 1996, p.83), where a ‘process’ is a time-ordered sequence of states a system takes in a given time period.<sup>3</sup>

Hartmann’s characterisation, which I take to be basically correct, opens some interesting questions. Consider my previous example: in order for the

---

<sup>1</sup> I am here referring standard maps on paper only. Giere (1999, pp.44-47) provides a detailed discussion of the function of maps as models.

<sup>2</sup> I am paraphrasing Hartmann here, for he speaks of models as if they were linguistic entities (made of ‘assumptions’, for example), whereas in this paper I follow the Semantic approach and take them to be objects.

<sup>3</sup> I am here modifying slightly Hartmann’s (1996) own definition of ‘process’ as “some object or system whose state changes in time” (p.83).

map-plus-miniatures to be a simulating device, the system must be capable of taking different states (the miniatures must change their position on the map, for instance). This means that there must be an agent prompting the changes in the system itself. Such a role may be played for instance by the officials in the army's head-quarters. Thus, counter-intuitively perhaps, the officials must *belong to* the simulating device itself. If 'simulation' is an anthropomorphic or more in general agent-dependent notion, as Bunge seems to suggest, we should not be troubled by these features. It is just natural that what is to be included and what to be excluded in a simulating system is partly arbitrary and/or dependent on one's interest. Simulations are not in nature, it is us who 'see' them and often build them according to our purposes. Similarly, a checkerboard and some pawns cannot by themselves simulate anything – although they can *represent* something: for example the state of a given battle at time *t*. A chequerboard, some pawns, *and two players* can simulate a battle or a war by representing a sequence of states of that battle or war. Most often, a simulating device will have some mechanism built into it, which once triggered will make the system go through a series of states automatically. The agent's role, then, will be merely that of setting the initial state and starting the process, which will then keep running until it is exogenously interrupted or runs out of steam.

### 3. SIMULATIONS VS. EXPERIMENTS: THE EPISTEMIC ACCOUNT

The distinction between simulations and experiments is more tricky than the one between models and simulations. In everyday scientific talk, such a distinction is certainly loaded with epistemic connotations: simulations are supposed to be somehow less fertile than genuine experiments for the production of scientific knowledge. Their results are often qualified as 'mere' simulations not to be mistaken for the 'real thing' (i.e. the real-world system whose behaviour is being simulated, or an experiment on the real-world system). The interesting question, however, is whether the epistemic difference is fundamental from a semantic viewpoint, or whether it is just a by-product of some more basic difference between experiments and simulations.

I should make clear that I am not interested in conceptual distinctions *per se*. My primary aim is to make sense of certain 'ambiguous' tools that are widely used in science. This is not a purely philosophical concern: scientists worry about the same issues – probably even more than philosophers do. Take for example the sort of laboratory work done by psychologists and economists interested in behavioural decision making. The

psychologist Baruch Fischhoff represents practitioners' worries by means of a graphic example. In the psychology lab, choices look like this:

**Choice A** In this task, you will be asked to choose between a certain loss and a gamble that exposes you to some chance of loss. Specifically, you must choose either: Situation A. One chance in 4 to lose \$200 (and 3 chances in 4 to lose nothing). OR Situation B. A certain loss of \$50. Of course, you'd probably prefer not to be in either of these situations, but, if forced to either play the gamble (A) or accept the certain loss (B), which would you prefer to do? (Fischhoff, 1996, p.232)

But in the real world, choices look like *this*:

**Choice B** My cousins ... ordinarily, I'm like really close with my cousins and everything. My cousin was having this big graduation party, but my friend – she used to live here and we went to ... like started preschool together, you know. And then in 7th grade her stepdad got a job in Ohio, so she had to move there. So she was in Ohio and she invited me up for a weekend. And I've always had so much fun when I'd go up there for a weekend. But, it was like my cousin's graduation party was then, too – like on the same weekend. And I was just like I wanted to go to like both things so bad, you know. I think I wanted to go more to like up Ohio, you know, to have this great time and everything, but I knew my cousin – I mean, it would be kind of rude to say, "Well, my friend invited me up, you know for the weekend." And my cousins from out of town were coming in and everything. So I didn't know what to do. And I wanted mom to say, "Well, you have to stay home", so then I wouldn't have to make the decision. But she said "I'm not going to tell you, you have to stay home. You decide what to do." And I hate when she does that because it's just so much easier if she just tells you what you have to do. So I decided to stay home basically because I would feel really stupid and rude telling my cousin, well, I'm not going to be there. And I did have a really good time at her graduation party, but I was kind of thinking I could be in Ohio right now. (Fischhoff, 1996, p. 232)

What do choices in environments like the first one tell us about behaviour in environments like the latter? And what *are* environments like the former anyway? Are they *simulations* of real-life situations, or are they *experiments* on human decision-making? One possible answer is that experiments like 'Choice A' test subjects' 'pure' cognitive capacities. But this is unsatisfactory: decision processes may be completely different in the two circumstances, and 'purity' is a poor consolation if it is unlike anything we are ultimately interested in explaining and understanding. A more reasonable answer is that in situations like 'Choice B' there is just too much going on,

and simplified settings like ‘Choice A’ are just intermediary steps on the way towards the understanding of complicated ‘real-world’ decision making. Indeed, it is always useful to think of experimental work as involving (at least) three different systems and two distinct hypotheses (see Figure 1).

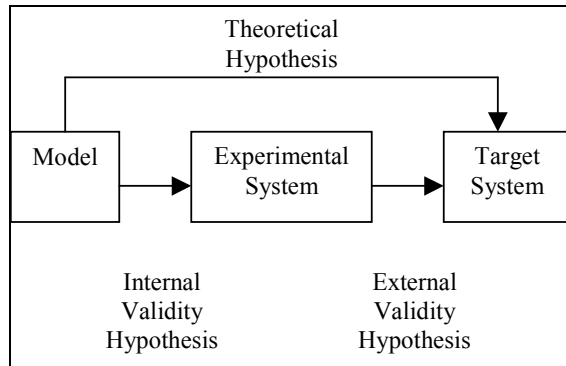


Figure 1. From models to target systems

As semantic theorists point out, models must be put in correspondence with their ‘real-world’ target by means of a theoretical hypothesis specifying what sort of relation (isomorphism, analogy, similarity, etc.) is supposed to hold between them. But the path between a model and the real world can be split into a number of sub-steps. Most often, the relation between a model and an appropriately designed experimental system is tested first (in the lab). Then it is further hypothesised that the experimental system stands in some specific relation with the target. This is typical of sciences, like medicine or economics, which are ultimately aimed at explaining phenomena in the unconstrained world outside the lab. Since the quality of field data is usually too poor to answer specific questions concerning causal mechanisms and processes, physiologists and social scientists find it useful to try and answer them first in the laboratory. Then, they show (or argue, or assume) that the laboratory was a correct replica of the real-world systems they were interested in the first place.

What kind of correspondence relation is specified by the external validity hypothesis? Given the amount of controversy concerning the nature of theoretical hypotheses in the semantic approach, one may be pessimistic about the prospects of finding a general answer to this question. But, minimally, an external validity hypothesis should map elements of the experimental system onto elements of the target system. These elements should be entities, properties, and relations between these entities and properties. (Causal relations should be prominent, one would think.)

Experimental psychologists like Fischhoff are worried because the fact that X causes Y in circumstances A (the lab), does not imply that X is a cause of Y also in circumstances B (the target). An external validity hypothesis or assumption bears the weight of any inference from the social science laboratory to the real social world.

‘Model’ and ‘simulation’ talk is more common in sciences, like economics or medicine, in which external validity worries are widespread and taken seriously. Indeed, you do not find in physics the sort of *a priori* scepticism towards experiments that you find, for instance, among economists. But the question is: should the distinction between experiments and other ‘mediating’ devices be based on such epistemic differences? At least two arguments can be levelled against this solution. Firstly, epistemic degrees of confidence in a particular scientific tool or device change in time. Thus, Galileo’s experiments were not immediately greeted with enthusiasm, but it would seem odd to claim that for this reason they should have been labelled as ‘simulations’ until they were accepted by the scientific community. What a scientific device *is* should not depend on whether we are confident in the reliability of its results.<sup>4</sup>

Secondly, simulations require an external validity hypothesis too, which may or may not be true depending on the circumstances. If I simulate the battle of Waterloo using little soldiers and horses, I work on the hypothetical assumption that, for example, the speed of horse miniatures stands approximately in the same relation with the speed of little soldiers on my map as the speed of infantry units stood with the speed of cavalry units in 1815. Only under this hypothesis can I use the simulating model to investigate, for instance, what would have happened if Napoleon had chosen a different strategy. The difference must lie elsewhere. Perhaps in the fact that psychologists and economists use human beings (students, for example) as subjects in their experiments. The intuition is that, unlike pawns and armies, toys and troops, human beings are *the same* in and out of the lab. But how exactly?

---

<sup>4</sup> Notice that I am not claiming that the difference between experiments, models and simulations is not anthropocentric. Indeed, as I shall clarify later, the distinction depends very much on the way we use these devices. The point here is merely that the distinction does not depend on the *degree of confidence* we attach to the correspondence relation between mediating and target system.

#### 4. SIMULATIONS VS. EXPERIMENTS: THE ONTOLOGICAL ACCOUNT

A material model of the propagation of light, according to the wave theory, can be built with the aid of water in a ripple tank. At a general level of analysis any kind of wave can be modelled as a perturbation in a medium determined by two forces: the external force producing the perturbation, and the reacting force working to restore the medium at rest. General relationships such as Hooke's law or D'Alembert's equation may hold for *all* kind of waves. More fundamental relationships, such as Maxwell's equations, describe the properties of the electric and the magnetic field only. The values given by Maxwell's equations can be used in D'Alembert's wave equation in order to obtain, for instance, the velocity of propagation of an electromagnetic wave, because electricity behaves *like* a wave, although the fundamental principles at work are different from those at work in case of, e.g., water waves. The terms appearing in the equation describing the target and the model-systems are to be interpreted differently in the two cases: the forces are different in nature, and so are the two media in which waves travel. The similarity between the theoretical model of light waves and the ripple-tank model holds at a very abstract level only. The two systems are made of different 'stuff': water waves are not light waves. Because of the formal similarity, though, the behaviour of light waves can be *simulated* in a ripple tank. Both light waves and water waves obey the same non-structural law, despite their being made of different 'stuff'. This is due to different reasons in each case: different underlying processes produce similar behaviour at an abstract level of analysis.<sup>5</sup> Similarly, human behaviour can to a certain extent be simulated by means of computerised models, but arises from 'machines' made of flesh, blood, neurons, etc. rather than silicon chips.

Herbert Simon (1969, pp.15-18) puts it as follows: simulations rely on a process of abstraction from the fundamental principles governing the behaviour of the simulating and the target systems. If similar 'organisational properties' arise at a given non-fundamental level from different substrata, it is possible to abstract from the substrata and simulate the behaviour of a system A by observing the behaviour of another system B which happens to

---

<sup>5</sup>Of course, if one believes in the reductionist story according to which everything physical is made of the same fundamental sub-atomic particles, then both light and water waves are 'made of the same stuff'. But the reductionist story is controversial (photons seem to have different properties from other particles), and at any rate the fact that everything is made of the same stuff does not play any relevant role in explaining why both systems display certain non-fundamental relations.



(or which is purposely built so as to) display those non-fundamental properties. Working on this idea, we can devise a criterion to demarcate genuine experiments from ‘mere’ simulations. The difference lies in the kind of relationship existing between, on the one hand, an *experimental* and its *target* system, and, on the other, a *simulating* and its target system. In the former case, the correspondence holds at a ‘deep’, ‘material’ level, whereas in the latter the similarity is admittedly only ‘abstract’ and ‘formal’. It is tempting to claim that in a simulating device the simulated properties, relations, or processes are generated by *different* (kinds of) *causes* altogether. Such a claim is problematic if you endorse a formalistic view of (type-level) causation – for example a view that defines causation in terms of purely probabilistic relations. If, in contrast, one takes causation to be a substantive property of specific kinds of systems (à la Wesley Salmon, for instance)<sup>6</sup>, the problem may disappear. In a genuine experiment the same ‘material’ causes as those in the target system are at work; in a simulation they are not, and the correspondence relation (of similarity or analogy) is purely formal in character.<sup>7</sup>

## 5. THE METHODOLOGY OF ‘PURE’ SIMULATIONS

To sum up: according to Simon, the difference between experiments and simulations has nothing to do with the (lack of) confidence attached to the correspondence relation between the simulation/experiment and its target, but rather with its very nature. The mapping in a simulation is purely formal, whereas experiments are materially analogous to their target systems. Because of the different nature of the correspondence relation, simulations and experiments are appropriate research tools in different contexts. Typically, simulations are used in one of two different ways: either (1) to bootstrap from the fact that a given effect (which we have observed in system A) can be produced by means of simulation B, to the fact that the relations governing the behaviour of B also govern the behaviour of A. Or (2) to argue that a certain effect observed by simulating with B will also be observed in the case of A because the two are governed by similar relations.<sup>8</sup> Both procedures are knowledge-producing ones. The point to be stressed is

---

<sup>6</sup> The latest versions of Wesley Salmon’s theory of causation, however, are of little help outside the realm of physics.

<sup>7</sup> See Mary Hesse’s (1963) distinction between ‘material’ and ‘formal’ analogies (p.63f).

<sup>8</sup>One may be unable to experiment with A, or the equations describing A may be so complicated that they can be solved only by means of some ‘brute-force’ solution in B.

that in both cases the governing relationships have to be fully specified for the simulations to be carried on. Systems of this kind are ‘transparent boxes’, to which the old *dictum* applies: “a simulation is no better than the assumptions built into it” (Simon, 1969, p.18).<sup>9</sup>

Geologists working on stratigraphy, for instance, study the structure of rock layers below the earth’s surface. They also investigate the process of strata formation, but have to face very serious obstacles, such as the impossibility of doing controlled experiments (processes of sedimentation last for millennia, and of course the target systems are too large to be manageable in a lab), the difficulty to gather data even about the present geography of the strata, the strong theory ladenness of the interpreted data, the complex interdependencies within geological systems, and so on. In order to solve at least *some* of these problems, geologists have devised simulation techniques like STRATAGEM, a computer-based modelling package used by large companies such as Shell Oil.<sup>10</sup> This simulation device works on the basis of a number of structural equations taken from the theory of ‘sequence stratigraphy’. The equations model the system’s outcome (the actual sedimentation) as a function of a number of variables including the hydrodynamics of sediment deposition, the subsidence patterns, the global sea level, the amount of sediment supplied to the basin, etc. The outcome of the simulation is dependent on the approximate validity of the theory of sequence stratigraphy, and also on the correct specification of the initial conditions and of the values assigned to the free parameters in the equations (incidentally, these are all problematic assumptions to be made in the specific case). Geologists try to simulate systems A (real-world geological structures) by means of a computer-model B, and all the fundamental relations in B must be known and specified in advance.

In figure 2, and in the diagrams that follow, I represent the presently unknown features of a system by means of dotted lines.<sup>11</sup> The question mark denotes an aspect of a system whose nature or functioning the scientists is investigating in a given case.

---

<sup>9</sup> Which does not mean that they are *just* as good as that: to run a ‘good’ simulation involves the use of approximations, computer implementation, etc. and thus requires more knowledge and skills than simply specifying the correct basic equations.

<sup>10</sup> I have learned about simulation techniques in geology from Francis Longworth and his unpublished paper on the methodology of STRATAGEM.

<sup>11</sup> In fact, a more precise representation would involve at least *two* systems (say, **A** and **B**), each one made of entities/properties and relations (say,  $X_A \rightarrow Y_A$ ,  $X_B \rightarrow Y_B$ ), plus some ‘external validity’ hypothesis stating that the two systems stand in a certain relation to each other. The dotted lines in my figures are shorthands for cases in which such a hypothesis is *not* known or well-confirmed.




Figure 2. Pure simulations in geology

Notice that some features of simulating systems like STRATAGEM is exactly opposite from that envisaged in the epistemic account. In order for these simulations to be successful, geologists must be very confident that the (formal) correspondence relation between relations in the simulating device and relations in the target holds to a certain degree of approximation. Experiments, as we shall see in the next section, do not require that much.

## 6. THE METHODOLOGY OF ‘PURE’ EXPERIMENTS

The crucial presumption behind experiments is that relevant components of the laboratory system are made of the same ‘stuff’ as those of the target system. (We shall see in a moment what ‘relevant’ means.) The experiment should feature the same causal processes that are at work in the real world, rather than just display some formal relation by means of a device made of different ‘stuff’. Experiments are particularly useful when one has an imperfect understanding of some basic causal mechanism of the system under study. They can be used in these contexts because the laboratory ‘stuff’ is assumed to be the same as the non-laboratory ‘stuff’.

What is unknown is often (but not always) what is under test. An experiment can give us more confidence in a theoretical model, if the theory makes some contestable assumption about some component of the target system, and if the experiment includes the real component (for example real human behaviour, as in experimental psychology and experimental economics). An experiment that merely reproduces all the assumptions of the model, for example by paying subjects to act according to the behavioural theory of the model, does not test anything at all (except perhaps the incentive system). But notice that not all that is imperfectly understood needs to be under test. For instance, one can test the efficacy of a drug without a detailed understanding of the epidemiological characteristics of a disease. The efficacy of the drug, rather than the process of infection, is what

is under test. Or you can do experiments on market behaviour even without a proper understanding of the mechanisms of individual choice and belief formation. Market institutions, instead of individual behaviour, are under test in these experiments. Subjects may trade at a certain equilibrium price because they are acting in a fully rational way, or perhaps because they are following some rule of thumb, or even by sheer imitation. Whatever the real causal process, we can use laboratory tests to study selected aspects of specific real-world economies as long as we are confident that the same (unknown) basic principles of behaviour apply in both cases.

Considerations of this kind have sometimes been proposed by experimentalists to defend their methodology. The economist Vernon Smith (1991), for example, argues that “the laboratory becomes a place where real people earn real money for making real decisions about abstract claims that are just as ‘real’ as a share of General Motors”. For this reason, “Laboratory experience suggests that all the characteristics of ‘real world’ behavior that we consider to be of primitive importance [...] arise naturally, indeed inevitably, in experimental settings” (pp.100-1). This reasoning supports experimenters’ confidence in their results. To them, the ‘real’ character of experimental markets helps to bridge the gap between a theory and its intended target of application. “Laboratory microeconomies are real live economic systems, which are certainly richer, behaviorally, than the systems parametrized in our theories” (pp.254-5). Experimental economies are indeed supposed to work according to the same principles as the target systems in the intended domain of economic theory, because the *relevant* components of the laboratory system are made of the same ‘stuff’.

Thus, both experiments and simulations are knowledge-producing devices. But the knowledge needed to run a good simulation *is not quite the same* as the one needed to run a good experiment. When reproducing a real-world system in the laboratory, the relationships describing the behaviour of both systems may not be known in advance. But one does not have to specify the full set of structural equations governing the target system. The trick is to *make sure* that the target and the experimental system are similar in most relevant respects, so as to be able to generalise the observed results from the laboratory to the outside world. Experimenters make sure that this is the case by using materials that resemble as closely as possible those of which the parts of the target system are made. They also make sure that the components of the ‘mimetic’ device are put together just like those of the target, and that nothing else is interfering. Of course, quite a lot of knowledge is required in order to do so, but no fundamental theory of how the target system works is needed. Parts of the laboratory system can be put between brackets and used as ‘black boxes’. The same processes, the same causal principles are supposed to be at work in both cases. Experimental

systems are reliable if they are made of the same ‘stuff’ as real world economies. No process of abstraction from the material forces at work is needed in order to draw the analogy from the laboratory to the outside world. One may abstract from ‘negligible’ causal factors, but not from the basic processes at work. The similarity is not merely *formal*, but holds at the *material* level as well.

## 7. HYBRIDS: SIMULATING EXPERIMENTS, EXPERIMENTAL SIMULATIONS

The distinction between simulations and experiments borrowed from Simon seems to be of the ‘black-or-white’ sort. Either the ‘stuff’ is the same as that of the target system, or it is not. But it is more realistic to allow for intermediate cases between the two extremes. This can be done in at least two ways. Mary Morgan (in this volume) discusses ‘hybrid’ entities that are neither entirely simulations nor entirely experiments, but a little bit of both. She focuses on cases in which some ‘materiality’ is transferred from the target system to the mediating entity, and reproduced therein, only to a certain extent. These hybrids are ‘quasi-material’ entities, as she puts it. The second strategy, in contrast, which I would like to explore in this paper, is to look at hybrids that combine purely experimental and purely simulating components. In fact, if simulations and experiments produce novel scientific knowledge in different ways, they must be partly complementary, and we should be able to combine them in the same project to exploit the potential of both.

My example comes once again from the social sciences. Experimental psychologists and economists are often concerned with designing experiments that reproduce in all ‘relevant’ respect real-world decision situations. Subjects, for example, are invited to trade goods in an environment governed by the rules of some real-world institution (say, a certain auction system). Even where realism is sought, however, experimenters sometimes have to make use of ‘artificial’ devices. Take, for example, experiments about so-called ‘common value’ goods – items whose value is the same, but unknown, to all traders. In experiments of this kind uncertainty is customarily implemented by means of a random draw. The subjects trade lottery tickets, in other words, which will be played out at the end of the experiment. Here uncertainty is *simulated* by means of a random draw. Uncertainty arises from the interaction of experimental subjects with a random lottery device, rather than with a ‘real’ good of unknown value (say, a concession for an oil tract or a licence for mobile phone frequencies). Before you use such a device, you need to be very confident that such a

way of modeling *that* particular aspect of the target system is legitimate. One has to be reasonably sure, in other words, that *that* part of the theory is right. Here more ‘paradigmatic’ sciences like economics are seemingly (but misleadingly, perhaps) better off: economists are confident to take such a move because auction theory prescribes to model agents’ uncertainty as a probability distribution of this sort. If challenged, they reply (as Robert Sugden, for example, has done in private correspondence) that the two phenomena (uncertainty faced by oil companies bidding for real tracts, and uncertainty faced by subjects in a laboratory auction) are particular instances of the same phenomenon – *uncertainty*, full stop. To ask whether uncertainty has been produced correctly in the laboratory would be like asking whether in an experiment to investigate the properties of steam the boiler has been heated by burning coil rather than by a Bunsen burner: heat is heat however it is produced, and obeys to the laws of thermodynamics.<sup>12</sup> Psychologists, in contrast, tend to be suspicious of assumptions of this sort and argue that human behaviour (people’s reactions to the same stimuli) vary widely from context to context. Of course there may be mechanisms general enough to support inferences to a wide range of circumstances, but that is an empirical matter to be settled *after*, not *before* experimentation takes place.




Figure 3. Experimental simulations

The structure of such ‘experimental simulations’ may be represented as in Figure 3. The hypothesis under test here is signalled by the question mark. Other causal factors are part of the target system, but must not necessarily be reproduced identically in the experimental design. For the experiment to work, the effect of the environment may be simulated by a device that ‘does the same thing’ as the corresponding real-world factor. But for that to be the case, we must *know* that the relation between ‘simulating’ and ‘simulated’ factors is one of formal similarity. Or, if you are not confident that the simulated features *really* closely resemble those of the target system, you

<sup>12</sup> I owe this example to Bob Sugden.

must make sure that they falsify reality in the ‘right’ way. If you want to test the cognitive capacities of human beings, for example, you may want to create an environment that is much less ambiguous and indeed tidier than the real world. The reasoning being that if subjects cannot read the situation correctly there, surely they will not be able to do it in messier circumstances. Sometimes ‘hyper-realism’ is better than no realism or even tentative realism.

## 8. CONCLUSION

Experiments, simulations, and models belong to the same category of scientific tools. They all are somehow ‘in between’ what we say about the world and the world itself. When we use them, we must be careful that they stand in the right relation both to what we say, and to the part of the world we are interested in. This is quite different from the standard account, where models and simulations stand on the theory’s side at one end of the spectrum, and experiments stand on the world’s side at the other end. It is also different from the standard view of testing, according to which scientists ask very general questions about the world, design very tight and controlled experiments to answer such questions, and the data will hit the theory no matter what. For the data to hit the target, the experiment, simulation, or ‘simulating experiment’ must mirror the target in a number of relevant aspects. Whether the mirroring should be purely formal or material in character, depends on the kind of question we are asking, and on the amount and quality of background knowledge we have accumulated regarding the target system itself. Take Fischhoff’s ‘Choice A’: is it an experiment or a simulation? *A priori* there is no answer to such a question. It depends what the scientists were aiming at in the first place, what sort of target systems they had in mind, which *aspects* of the target they were investigating, and crucially what they already knew about the relationship between their targets and ‘experiments’ (or ‘simulations’) like ‘Choice A’.

## ACKNOWLEDGEMENTS

Many helpful criticisms and suggestions were provided by Dan Hausman, Mary Morgan, Bob Sugden, John Worrall, and members of the audience at MBR01. All the remaining mistakes are mine.

**REFERENCES:**

- Bunge, M., 1969, Analogy, simulation, representation, *Revue internationale de philosophie* 87:16-33.
- Fischhoff, B., 1996, The real world: what good is it?, *Organizational Behavior and Human Decision Processes* 65:232-248.
- Giere, R., 1999, Using models to represent reality, in *Model-Based Reasoning in Scientific Discovery*, L. Magnani, N.J. Nersessian, and P. Thagard, eds., Kluwer, New York, pp.41-57.
- Guala, F., 1998, Experiments as mediators in the non-laboratory sciences, *Philosophica* 62: 901-918.
- Hartmann, S., 1996, The world as a process: simulations in the natural and social sciences, in *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, R. Hegselmann, U. Mueller, and K. Troitzsch, eds., Kluwer, Dordrecht, pp.77-100.
- Hesse, M.B., 1963, *Models and Analogies in Science*, Sheed & Ward, London.
- Longworth, F., 1998, Computer modelling in the earth sciences, *Imperial College of Science, Technology and Medicine working paper*.
- Morgan, M.S., 2001, Model experiments and models in experiments, *this volume*.
- Morrison, M. and Morgan, M.S., 1999, Models as mediating instruments, in *Models as Mediators*, M.S. Morgan, and M. Morrison, eds., Cambridge University Press, Cambridge, pp.10-37.
- Simon, H.A., 1969, *The Sciences of the Artificial*, MIT Press, Boston.
- Smith, V.L., 1991, *Papers in Experimental Economics*, Cambridge University Press, Cambridge.