Measurement in Physics and Economics Project Discussion Paper Series,
DP MEAS 25/02, London School of Economics 2002.

Models and Representation:

Why Structures Are Not Enough.

Roman Frigg

1. Introduction

Models occupy a central role in the scientific endeavour. Among the many purposes they serve, representation is of great importance. Many models are representations of something else; they stand for, depict, or imitate a selected part of the external world (often referred to as target system, parent system, original, or prototype). Well-known examples include the model of the solar system, the billiard ball model of a gas, the Bohr model of the atom, the Gaussian-chain model of a polymer, the MIT bag model of quark confinement, the Lorenz model of the atmosphere, the Lotka-Volterra model of the predator-prey interaction, or the hydraulic model of an economy, to mention just a few. All these models represent their target systems (or selected parts of them) in one way or another.[1]

[1] I do not hereby suggest that *all* models represent, some may merely be tools for manipulating or intervening in the world.

This is essential to their functioning in a context of investigation. In order to be a source of knowledge a model must be representative. A model can tell us about the nature of reality only if we assume that some aspects of the model have counterparts in the world. Hence, if we wish to learn from a model about the world – and I take it that this is often the case – we are committed to the claim that the model involves some sort of representation.

This raises the question of how models manage to represent and of what kind of things they are. The currently most influential answer to this question has been given within the context of the so-called semantic (or model-theoretic) view of theories. At the core of this view lies the notion that models are structures (where a structure, roughly speaking, is a collection of objects along with the relations in which they enter) and that they represent due their being isomorphic to their target system. For this reason I shall refer to it as the *structuralist view of models.* (This coincides with the terminology used by its advocates. While the term 'semantic view of theories' has been widely used in the past, more recently some of its proponents – notably van Fraassen, DaCosta, and French – have decided to call themselves 'structuralists'. I use these terms interchangeably in this essay).

The aim of this paper is to argue that this is a dead end. Despite allegations to the contrary, the structuralist conception cannot account for the fact that models represent. I shall argue for this by undermining the structuralist conception 'from within', rather than putting forward criticism 'from outside' (such as that it cannot account for scientific practice or theory change, to mention just two frequent objections). That is, I take the structuralist conception of models at face value and evaluate how well it conforms with its own standards; more specifically, I discuss the question whether the structuralist view can account for the fact that models represent (something that structuralists accept) and I conclude that it cannot. In a nutshell, my argument runs as follows: Trivially, structures *per se* do not represent anything at all. So we have to add some 'stuff' in order to endow them with representative power. But what is this stuff? A rather lengthy discussion leads me to the conclusion that the 'stuff' we have to add is a *physical design,* by which I mean roughly something like a model in the straightforward scientific sense (which is closely related to Cartwright's prepared descriptions, Achinstein's theoretical models, Suppes' physical models, and Hesse's analogical models). However, if structures cannot be endowed with representative power without recourse to a physical design, structures by themselves are not models. Models are representative devices and everything that is essential to this purpose is part of the model. For this reason I conclude that a model is a complex entity consisting (at least) of a structure, a physical design and a process that hooks up the two. But this is not structuralism any more. Summing up, we have a kind of reductio of the structuralist position: In trying to make structuralism work one is forced to tack on elements which structuralism does not allow for.

The thrust of this argument is to link up a philosophical account of models with a more straightforward scientific usage of the concept. This is not a new project, it can be traced back (at least) to Hesse (1953, 1963) and Achinstein (1965, 1968). However, my claim is stronger than theirs. I argue in this paper that physical designs are an essential part of any representational model, and not merely a heuristic tool for working scientists or a 'friendly' but in essence unnecessary amendment. For short, my claim is: No physical design, no representation.

I should conclude this introduction with a disclaimer: I do *not* put forward the view that structures are unimportant, uninteresting, or what have you – structures are extremely important in any mathematised science, they are by themselves just not enough to make up a representational model.

2. The Structuralist View of Models in Science

In this section I shall introduce the structuralist account of scientific models. But before doing so it seems worthwhile to say a few words about structures: A structure S is a composite entity consisting of the following ingredients: (i) a non-empty set U of individuals called the domain (or universe) of the structure S, (ii) a set O of operations on U (which may be empty), and (iii) a non-empty set R of relations on U (Machover 1996, 149; compare Bell and Machover 1977, 9, 49, 162; Hodges 1997, 2; Boolos and Jeffrey 1989, 98-9). Often it is convenient to write these as a ordered triple: S=<U, O, R>. Note that nothing about what the objects are matters for the definition of a structure – they may be whatever one likes them to be. Similarly, operations and functions are specified purely extensionally; that is, n-place relations are defined as classes of n-tuples, and functions taking n arguments are defined as classes of n+1-tuples.

The crucial move now is to claim that scientific models are nothing but structures in this sense. This view is nicely encapsulated in Suppes' by now famous slogan that 'the meaning of the concept of model is the same in mathematics and the empirical sciences.' (1960a, 12; compare 1960b, 24; 1970, Ch.2 pp. 6, 9, 13, 29). This has become the cornerstone of the semantic view of theories. Although there are differences in detail between the different proponents of this approach, it is the shared belief that the relevant notion of model for the empirical sciences is the one we find in mathematical logic (or one very closely related to it).[2]

[2] The precise formulation of what these models are varies from author to author. A survey of the different positions can be found in Suppe (1989, 3-37). How these accounts differ from one another is an interesting issue, but for the present purposes nothing hinges on it since, as Da Costa and French (2000, 119) correctly remark, '[i]t is important to recall that at the heart of this approach [i.e. the semantic approach as advocated by van Fraassen, Giere, Hughes, Lloyd, Thompson, and Suppe] lies the fundamental point that theories [construed as families of models] are to be regarded as *structures*.' (original emphasis)

However, in themselves structures are not representations of anything in the world. They are pieces of pure mathematics, devoid of empirical content. A representation must posses 'semantic content', that is, it must stand for something else. But structures *per se* do not stand for anything at all. They do not indicate any real-world system as their object.

Structuralists are well aware of this problem and acknowledge that structures, in order to be representative models, must be supplemented with a specification of the relation they bear to the target system. What is the nature of this relation? In keeping faithful to the spirit of the semantic view, the most natural choice is structural isomorphism: A structure *S* represents a target system *T* iff they are structurally isomorphic. It is this isomorphism that links the structure to the world; it is due to their being isomorphic to some part of the world that structures get endowed with representative power and 'are about' something.[3]

This is in need of some qualification. What does it mean for a target system to be isomorphic to a structure? After all, isomorphism is a relation that holds between two structures and not between a structure and a piece of the real world. The essential (and implicit) assumption in the isomorphism claim is that the target system exhibits a certain structure. Target systems are not 'bare things', they are structured in a certain way; in other words, they exhibit a 'physical structure' $S_p = \langle U_p, O_p, R_p \rangle$. This is an important and by no means trivial assumption. I shall discuss it in great detail later on, but for the time being I assume that it is correct and unproblematic. Given this, it is now possible to make the concept of an isomorphism precise. An isomorphism is a mapping $f: S_p \rightarrow S$ such that (i) f is one-to-one (bijective), (ii) f preserves the system of relations in the following sense: If the elements $a_1, ..., a_n$ of S_p satisfy the relation R^p then the corresponding elements $b_1=f(a_1), ..., b_n=f(a_n)$ in S satisfy R, where R is the relation in S

[3] This is the view we find in van Fraassen (1980, Ch. 2; 1989, Ch. 9), French and Ladyman (1999), French and Da Costa (1990), French (2000), and Bueno (1997 and 1999), to mention just a few.

corresponding to R^p. And similarly, for all operations g^p of S_p we have $f[g^p(a_1, ... a_n)] = g(f(a_1), ..., f(a_n))$ where g is the operation in S corresponding to g^p.

Note that the notion if isomorphism as introduced here is symmetrical and reflexive, and transitive: If A is isomorphic to B, then B is isomorphic to A, every structure is isomorphic to itself, and if A is isomorphic to and B to C, then A is isomorphic to C.

To sum up, the structuralist view of models is the following: **A Model *M* is a structure; and *M* represents a target system *T* iff *T* is structurally isomorphic to *M*.**[4]

Against this view Mauricio Suárez (1999, 79) raised the following objection: The capacity of a model to represent must be an inherent part of it and not something that is added to it as an 'external factor'. More specifically, since structures *per se* do not represent anything in the world, it is inappropriate to refer to them as models. Models inherently 'point' to their targets and do not need to be connected with them by postulating a relation (structural isomorphism) external to the model to hold: '[...] a model is a representation, as it essentially intended for some phenomenon; its intended use is not an external relation that we can choose to add to the model, but an essential part of the model itself.'

I agree with this point of view. Taking a model to be a structure is just like taking a picture to be (merely) a collection of colour splotches on a canvas. However, proponents of the semantic view could now counter that this argument rests on a confusion about the use of the word 'model'. They may say that when they speak of models as structures what they actually

[4] I should mention that though this notion of a model is usually employed in the context of the semantic view of theories, it is not necessarily bound to it. As Da Costa and French (2000, 120-1) rightly point out, an account that emphasises the independence of models from theory (such as Morrison's or Cartwright's) could, in principle, still adopt a structuralist view of what models are and of how they relate to reality. Hence, nothing in what follows hinges on any particular feature of the semantic view of theories.

mean is ‘structure *plus* isomorphism’, it is just too cumbersome to repeat that every time. On this reading, the structuralist view is that a model is a set-theoretic structure *plus* the isomorphism that holds between the target and the structure – and if at times bare structures are called ‘models’ this is just sloppy talk.

Although I am not sure whether this route is open to all proponents of the view (van Fraassen, for instance, is reluctant to specify the intended applications of his models, see his 1980, 66), it seems to be a reasonable suggestion. Summing up we obtain: **The structure *S* represents the target system *T* iff *T* is structurally isomorphic to *S;* given that this is the case, the model *M* is the pair consisting of the structure *S* and the isomorphism that holds between *S* and *T*.**[5] Schematically we get the following picture:

[5] It is sometimes pointed out (Redhead 2001, 79) that not all elements of the structure have corresponding bits in reality. In this case the appropriate relationship between model and target system is embedding and not isomorphism. (We obtain the definition of an embedding if we stipulate that f, as specified in the above definition, is only injective and not necessarily one-to-one: For all x, y in A, if x is not equal to y, then f(x) is not equal to f(y).) For the purpose of the present discussion this difference does not matter for the following reason: If a structure R is embedded in a larger structure S then R is isomorphic to a substructure S’ of S; and it is the substructure S’ that has representative power. In this case the arguments to follow apply to S’ instead of S.

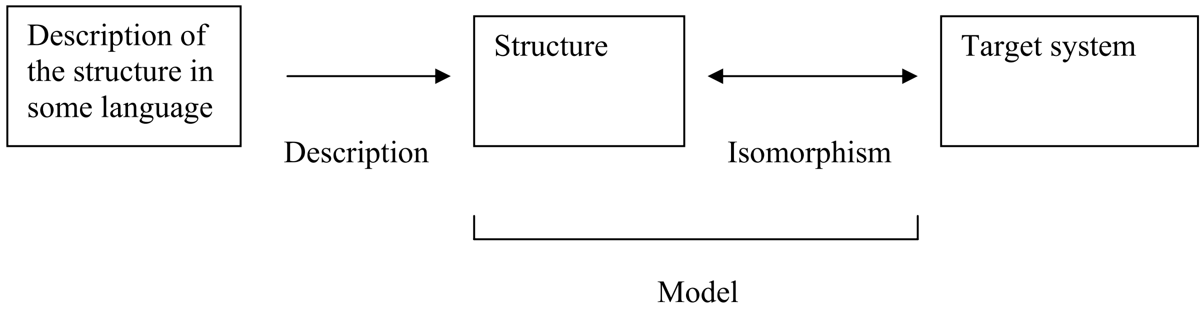

Fig. 1: The structuralist conception of representation.

3. Structure and Representation

Is the 'structure *plus* isomorphism' notion of models strong enough to assure that a model is a representation of something else, to guarantee that it possesses 'aboutness'? In this section I shall argue that this is not the case.

(1) Some science fiction to start with

Consider the following variation of a well-known thought experiment due to Hilary Putnam (1981, 1ff.):[6] Imagine a mathematician who is totally ignorant about physics. One day, sitting at the desk in his study, he writes down the equation $d^2x(t)/dt^2 = -k*x(t)$ and discusses its solutions with all sophistication and rigour of a modern mathematician. A person with some background in physics immediately realises that this is the equation describing the motion of an oscillating

[6] I should mention that Black (1970, 104) discusses a very similar thought experiment.

pendulum bob; or using the model-theoretic jargon: The phase space trajectory the equation describes is a *model* of the motion of the pendulum bob. But did the *mathematician* (who, recall, is totally ignorant about physics) write down a model of the moving pendulum?

On a little reflection most of us would probably say: no, *he* didn't. What the *mathematician* wrote down is a piece of pure mathematics, a mathematical structure if you like, but not a model of any physical system. The mathematician, after all, has never heard anything about oscillations, and he had no intention to come up with a model for an oscillating pendulum bob. He simply wrote down an equation and solved it. The fact that *someone else* can 'see' this structure as a model of a pendulum is not of his concern.

The point I want to get at is that structural isomorphism is not sufficient to make something represent something else. In Putnam's original thought experiment an ant is crawling on the beach leaving a trace that resembles Winston Churchill. Does this trace represent Churchill? He concludes that it does not:

> 'Similarity […] to the features of Winston Churchill is not sufficient to make something represent or refer to Churchill. Nor is it necessary.' (1981, 1)

> '[T]he ant's 'picture' of Winston Churchill has no necessary connection with Winston Churchill. The mere fact that the "picture" bears a resemblance to Churchill does not make it into a real picture, nor does it make it a representation of Churchill. Unless the ant is an intelligent ant (which it isn't) and knows about Churchill (which it doesn't), the curve it traced is not a picture or even a representation of anything.' (1981, 3)

Mutatis mutandis things are the same in the case of the ignorant mathematician. Unless he knows about oscillations and intends to come up with a model for them, his structures do not represent oscillations – regardless of whether they are isomorphic to oscillations or not.[7]

(2) Why structures don't represent

I have used Putnam's story to draw the attention to the problem of reference. But a story is not an argument yet. So before I discuss (possible) ways around the problem I have to provide further arguments to the conclusion that isomorphism is not enough to establish representation.[8]

(a) Structural Isomorphism has the wrong formal properties

The first (and somewhat simple) reason why representation cannot be cashed out in terms of structural isomorphism is that the latter has the wrong formal properties: Isomorphism is symmetric, reflexive and transitive while representation is not.[9]

---

[7] In passing I should note that there is a historical argument pulling in the same direction. Many mathematical structures have been discovered long before they have been used in science. Hilbert spaces, curvilinear geometry, matrix algebra, knot theory, and group theory, are points in case. If we subscribe to the view that mathematics refers 'by itself' we have to accept that Hilbert invented quantum mechanics, or that Riemann discovered general relativity. But this is obviously mistaken. It is only later that scientists started to *use* these mathematical tools to represent things and processes in nature.

[8] The arguments in this subsection I share in common with Mauricio Suárez with whom I have been discussing them over the last year.

[9] The first two points were earlier emphasised by Goodman (1976, 4) in connection with the similarity theory of representation.

Symmetry: If x is isomorphic to y then y is isomorphic to x (see above). But if x represents y, then y need not (and in fact in most cases does not) represent x. While a photograph may represent the Prime Minister, the Prime Minister does not represent the photograph; a tube map represents the London underground, but the underground does not represent the map. Likewise, Maxwell's 'billiard ball model' represents a gas, but not vice versa.

One might now be tempted to counter that the relation between the billiard balls and the gas, for instance, is actually asymmetrical because a scientist has picked out features of the balls to represent the gas molecules. This is a good point, and I shall come back to this suggestion later on, but it is not available at present. The view I am discussing seeks to cash out representation purely in terms of structural isomorphism and users have not yet entered the scene.

Reflexivity: Everything is isomorphic to itself, but most things do not represent themselves – I am isomorphic to myself, but I do not represent myself.

Transitivity: Isomorphism is transitive but representation is not. If a S_1 is isomorphic to S_2 and S_2 is isomorphic to S_3 then S_1 is isomorphic to S_3. This need not be the case for representation.[10] Take the bronze buste of Popper's in the corridor of the philosophy department. It represents Popper. If I take a photograph of the buste, the photograph represents the buste. But from this is it does not follow that the photograph represents Popper, and in fact it does not.

---

[10] There is a minor technicality to sort out at this point (thanks to Mauricio Suárez for drawing my attention to this). To claim that representation is not transitive may amount either to claiming that it is *non-transitive* or that it is *atransitive*. A relation R is non-transitive if '$\neg\forall x\forall y\forall z[(Rxy \,\&\, Ryz) \rightarrow Rxz]$' is true; that is, if there exist three individuals a, b, and c such that $Rab \,\&\, Rbc$ is true while Rac is false. A relation R is atransitive if '$\forall x\forall y\forall z\neg[(Rxy \,\&\, Ryz) \rightarrow Rxz]$' is true; that is, if there do not exist any individuals a, b, and c such that $(Rab \,\&\, Rbc) \rightarrow Rac$ holds. Non-transitivity is the weaker notion than atransitivity in that the latter implies the former but not vice versa. For the present argument the weaker notion will do: My claim is that representation is non-transitive.

An example from science for the non-transitivity of representation is the following: In chaos theory most of the by-now famous mappings such as the cat-map, the horse-shoe, the tent map, or the baker's transformation have not been designed to model any real process. They are simplified and schematic pictures of an exceedingly complex continuous phase flow (which is in fact too complicated to be studied directly) and it is this phase flow, which represents the dynamics of the real system. Hence, the mappings represent the phase flow and the phase flow represent the dynamics of the physical system, but the mapping does *not* represent the dynamics of the physical system.

(b) Structural Isomorphism is not sufficient for representation

Structural isomorphism is too inclusive a concept to account for representation. In many cases neither one of a pair of isomorphic objects represents the other. Two copies of the same book, for instance, are perfectly isomorphic to one another but neither is necessarily a representation of the other. Isomorphism between two items is not enough to establish the requisite relationship of representation; there are many cases of isomorphism where no representation is involved. Hence, isomorphism is not a sufficient condition for representation.

One might now argue that this critique is spurious since in the given set-up this problem cannot crop up at all. The models under consideration are structures and the target systems are objects in the world. Counterexamples of the aforementioned type can then be ruled out simply by building this matter of fact into the definition of representation: Introduce the ontological restriction that a model *M must* be a structure and the target *T must* be a concrete object (or process) in the world (that exemplifies a certain structure). This strategy obviously blocks the above objection. Even if one book resembles the other, it does not represent it because it belongs to the wrong ontological category.

Unfortunately things are not that easy. Though models often do refer to things in the world, this is not necessarily so. Just as a picture can represent another picture, a model can represent another model rather than anything in the world. Consider again the aforementioned case of chaos theory. Mappings like the baker's transformation are models that represent what happens in another model, not what is going on in the world. But this is incompatible with the amendment suggested above since by requiring that the target of a model must be a real system we would rule out such obvious and important cases of representation. This, however, is an unfortunate and unacceptable a result.

(c) Multiple realisability

Looking at successful applications of mathematics in the sciences we find that quite often the same structure is used several times in different contexts and even across different disciplines. Linear equations, for instance, are widely used in physics, economics, biology, and the mathmatised parts of psychology, to mention just a few. Similarly, ordinal measurement scales are used to quantify length, volume, temperature, pressure, electrical resistance, hardness of a solid and many other things. The $1/r^2$ law of Newtonian gravity is also the 'mathematical skeleton' of Coulomb's law of electrostatic attraction and the weakening of sound or light as a function of the distance to the source.[11] Harmonic oscillations are equally important in the context of classical mechanics and classical electrodynamics (for a detailed discussion of this case see Kroes 1989). These examples suggest that structures are 'one-over-many', as Shapiro (2000, 261) puts it. That is, the same structure can be exemplified by more than one target

[11] This fact has already puzzled Peirce (see Steiner 1998, 35).

system. Borrowing a term from the philosophy of mind, one can say that structures are *multiply realisable.*

And these are by no means isolated cases. Certain geometrical structures are exemplified by many different systems. Just think about how many spherical things we find in the world; or consider the spiral we know from helter-skelters and spiral staircases and that later on even came to serve as a model of DNA. The same is true of many mathematical theories as well. Take once more the case of chaos theory. Many of its mathematical structures have been put to use in various different fields such as physics, biology, economics, medicine, or climate research. And similar remarks apply to other branches of mathematics such as partial differential equations, for instance.

To see how the multiple realisability of structures clashes with the representative power of models, we need to bear the following feature of representations in mind. I observed at the outset that models are representations of *something else*. Implicit in this is that models are representations of some *particular* target system. The target can either be a token (as in the case of cosmological models) or a type (as in the case of models of the atom). But models are never models of several different things at the same time – one particular model is always a model of one particular thing. In that they are like pictures: A picture depicts a particular object, Big Ben say; but the very same canvass cannot at once be a picture of Big Ben, the Eiffel tower and the Campanile of San Marco. Representations, by their very nature, are directed towards one particular phenomenon; or to put it differently, they have the capacity to bring one particular thing to the mind of a suitably qualified audience.

It is this feature of representation that is incompatible with multiple realisability. If one particular set-theoretic structure is isomorphic to more than one system exemplifying the same empirical structure, which one is it a model of? What is the harmonic oscillator, to stick with the example, a model of? A pendulum bob? A lead ball on a spring moving up and down? The voltage in an electric circuit with a condenser and a solenoid? The amplitude of the B-field of an electromagnetic wave? The motion of atoms in the wall of a black body? It is isomorphic to all of them, but is it a *model of* all of them? I have just observed that one of the main characteristics of a model is that it is a model of one particular selected part of the external world. But if several parts of the external world can exemplify the same structure, a structure *per se* does not stand for anything in particular. For which one of all the structurally isomorphic systems does it stand? We face the dilemma that the structure *as a model* must stand for one particular system (or one particular type of system), but as a bare structure it is isomorphic to all of them and there is nothing in the set-up that picks out one of the systems as the 'privileged' one of which the structure really is a model. The structure fails to indicate to which one of the structurally isomorphic targets it should be applied. For this reason a structure (plus isomorphism) is not a model.[12]

(d) Identity conditions for models

A successful account of models has to give us identity conditions for them, enabling us to say under what conditions two models are the same. That is, we must be able to individuate models

[12] Van Fraassen mentions this problem at one point: 'This suggests that the intention, of which sorts of phenomena are to be embedded in what kinds of empirical substructures, be made part of the theory. I do not think so.' (1980, 66). He thinks that the problem of 'unintended interpretations' disappears as science progresses. I don't think that this is true. I cannot see any prima facie reason why the development of science should be such that one structure is used just in one particular context. On the contrary, it is good scientific method to build on analogies and to use structures one is familiar with to model new domains (Achinstein 1968, Ch. 7; Hesse 1966, Ch. 2; Psillos 1995, Sec. 3). But this suggests that the number of applications of one and the same structure is unlikely to decrease.

(recall the slogan: 'no entity without identity'). But if models are taken to be structures this is not possible. The argument takes the form of a *reductio ad absurdum* and runs as follows: Let A and B be two target systems, different from each other, that exemplify the same structure (which is perfectly possible as we have seen in the previous paragraph) and let M_A and M_B be the respective models (i.e. M_A is a model of A and M_B is a model of B). Since by assumption A and B are different (they may be a pendulum bob and a electric circuit, for instance), their models must be different as well since models are about a particular system. However, since A and B exemplify the same structure and since, by assumption, models *are* structures M_A and M_B must the same. So we end up with the contradiction that M_A and M_B are and are not identical. Hence one of the premises must be false and I take it that it is the one that models are (just) structures. This shows that structures are not enough to individuate a model.

(e) The possibility of inaccurate and false representation

Many representations are inaccurate in one respect or another. The Bohr model of the atom or the liquid drop model of the nucleus are well-known examples of models whose structure is not isomorphic to the structure of their respective target systems. But the structuralist view cannot take inaccurate models as representational since on this view either the model is representational and thus perfectly isomorphic to its target system, or it fails to represent altogether.

I take it that this is too restrictive. Many models are inaccurate in some way and nevertheless manage to represent their target systems – the two aforementioned examples are points in case. Therefore, structural isomorphism is not necessary for representation.

A further problem for the structuralist view arises in connection with models that are not only inaccurate but outright false. For thousands of years a flat disk has been used to represent the earth.[13] Though this is, as we now know, false the flat disk *is* a representation of the earth, albeit a wrong one. However, the structuralist view cannot regard false models as representational for the reasons mentioned above. Either the representation is accurate or it is not a representation at all; by definition, no representations can ever be false.

This view stands in stark contradiction to scientific practice. In order to assess the quality of a model we have to presuppose that it represents. Only when we assume that *M* represents *T* we can ask the question of how well it does its job. We tentatively put forward a model as a representation of something and then try to find out whether it is an accurate representation or not. But this becomes impossible if we deny that a false model has representative power. On what grounds could we say that the disk model of the earth is wrong if we deny its representational character? Therefore, to deny that false models represent undermines the process of testing a model and as a consequence makes research impossible.

(3) Where does representative power come from?

What lies at the heart of the above arguments is the problem of reference. Models refer to something, structures (even amended with isomorphism) do not. Hence the question is: In virtue of what does a model have reference? Can the structuralist account as formulated above be improved?

It seems that such an improvement is quite easy to get. What has been left out so far are possible users. Structures are not models 'in themselves', they become models when someone

[13] For those who prefer less colloquial examples, think of Thomson's model of the atom (now commonly, in a rather pejorative manner, referred to as the 'pudding model') or the fluid model of heat conduction.

uses them as such. That is, a model is representational because it is intended to so by a user. This is to say that representations are intentionally created; nothing may accidentally be a representation. Therefore users are an essential part of the picture. This suggests that as a way out of the difficulties we can invoke the intentions of the scientist who created the model.

From a formal point of view this amounts to making representation a triadic relation: a *user* takes *something* as a model of *something else*. Working this into the structuralist view of models we obtain the following: **The structure *S* represents the target system *T* iff *T* is structurally isomorphic to *S* and iff *S* is intended by a user to represent *T;* given that this is the case, the model *M* consists of the structure *S,* the isomorphism that holds between *S* and *T*, and the intention of the user.**

At first glance this appears to be a successful move, since in this version four of the above criticisms no longer go through: (a) The appeal to intentions renders representation non-symmetric, non-reflexive and non-transitive. So the problem with the wrong logical properties vanishes. (b) The problem that isomorphism is not sufficient for representation is resolved 'by definition' since we simply stipulate that S represents. (c) The problem of multiple realisability vanishes because the user can intend S to be a representation of some particular system (e.g. the voltage in an electric circuit) and forget about the other 'unintended' applications. (d) Building intentions into the definition of the model undercuts the contradiction because M_A and M_B are no longer the same. We are now left with objection (e), the argument from inaccurate and false representation. The present suggestion does not resolve this difficulty. However, let's assume for the sake of the argument that it is possible to get around these difficulties in one way or another (e.g. with a theory of idealisation).

Have we then solved the puzzle of representation? I don't think so. I agree that users are an essential part of any account of representation, but merely adding the condition that someone intends to use *S* as a model of *T* to the standard structuralist account is not enough. The problem with this suggestion is that it is incomplete. If we look at how the above-mentioned problems are resolved, we realise that it is the appeal to intention that does all the work and the original suggestion, the isomorphism between structure and target, does not play any role any more. This is unsatisfactory for (at least) two reasons: First, when we ask how representation works we would like to know what exactly a scientist does when she uses *S* to represent *T*. If we are then told that she intends to represent *T* by means of *S* this is merely a rephrasing of the problem and not an answer because what we want to know is what this intention involves. What exactly do we do when we intend to use *S* as a representation? To get the gist, consider the analogous problem in the philosophy of language: How do words acquire reference? We do not solve this problem merely by saying that a speaker intends words to refer to certain things. What we want to know is *how* the speaker achieves reference to something by using a word and coming to terms with this puzzle is what philosophers of language try to do in theories of reference (such as the description theory, the causal theory, the speech-act theory, etc.). The situation in the philosophy of science exactly parallels the one in the philosophy of language: What we have to understand is *how* the speaker achieves the use of *S* as a representation of *T* and to this end much more is needed than appeal to intention. In fact, what we need is a 'philosophy of science analogue' of theories of reference.

Second, a notion of representation based solely on intention is too liberal to account for *scientific* representation. Can I intend the teacup on my desk to represent the planetary system or can I simply put a dot on a piece paper and take it to represent an atom? As long as representative power solely rests on intention there is no way to rule out such cases – anything can be taken to represent anything, there are not limits. However, this is too loose a notion of representation to be of any use in science and this is certainly not what structuralists have in mind.

The situation can be improved if we try to bring the original idea that the model and the target must be isomorphic and the appeal to a user's intentions together. This does not seem to be too difficult. Notice that the only thing we really need in order to connect a structure to reality is the ability to identify individuals and relations in the world, to point to them and to say to which element in the structure they are supposed to correspond. More specifically, the process by which a structure S=<U, O, R> is endowed with representative power is something along the following lines: Identify an individual in the target system and match it with u_1. Then identify another individual and match it with u_2, and so on until all individuals in the target are matched with an individual of U; and similarly for the operations o_i and relations r_i of S. It is this job that intentions are supposed to do: Instead of taking the user's intention to connect S *as a whole* to T *as a whole* (as we did above), we now use them to associate basic parts of the structure (individuals, relations, functions) with parts of the target.

This suggestion is much less liberal than the previous one and effectively undercuts absurdities like the example with the teacup or the dot. However, there is still the other objection against intentions. I have observed that the appeal to intentions by itself is not enough to solve the problem of reference. But do I not presuppose that intentions actually do fix reference when I appeal to them in the process of matching elements of the structure with elements of the target since this matching amounts (mutatis mutandis) to the same as reference fixing? This is true, but there is a considerable difference between the two cases. The move from directing intentions towards a structure *as a whole* to directing them towards *parts* of a structure transforms an intractable problem into a tractable one. While it is rather mysterious how we could intend a structure as a whole to represent a target as a whole, we have firm intuitions about what it means for one element of the structure, u_2 say, to stand for an element of the target. In fact, the relation between the elements of the structure is roughly the same as the one between a name and its bearer. Once this is realised we can turn to the philosophy of language and use existing accounts of reference to cash out how this matching works (such as the causal theory of reference, for instance). – I should note that I do not claim that reference is an unproblematic concept, but even if all currently available accounts of reference suffer from certain difficulties, there is no doubt that reference as a phenomenon is real and that there must be some theory that cashes out how it works. Taking this for granted we can give the following statement of the amended structuralist view: **The structure *S* represents the target system *T* iff every element of *S* is intentionally matched with an element of *T* (by process similar to reference fixing in the case of language) and iff *T* is structurally isomorphic to *S*. Given this, the model *M* consists of the structure *S*, the isomorphism that holds between *S* and *T*, and the intentional matching.**

I take it that this suggestion, though never explicitly put forward by structuralists, captures quite well the underlying intuitions of a structuralist programme – at least there is nothing in this 'augmented' definition of representation that contradicts its spirit.

Where does this leave us? Are we now where we wanted to be? Were these amendments enough to remove the qualms over the structuralist account? Not quite. In all that has been said so far a non-trivial assumption has been made: That target systems exemplify structures. Recall, at the outset we assumed (as a working hypothesis) that the target system naturally exhibits an empirical structure $S_p=<U_p, O_p, R_p>$. This assumption is essential to what has been said so far. If the target is not structured (in the rigid set-theoretic sense) there is no meaningful way to claim that the target is isomorphic to the structure S because isomorphism is a concept that is applicable only to structures. Similarly, the dubbing process requires that there are identifiable and discernible individuals, relations and operations sitting there that can be matched with elements of the structure. If this is not the case, no dubbing can take place. For short, if we don't find the target system structured at the outset, the isomorphism cannot be set up and the required dubbing cannot take place. Therefore, the assumption that the target exhibits the empirical structure $S_p=<U_p, O_p, R_p>$ is absolutely vital.

Another way of putting the same point is the following: Whenever semanticists talk about how models relate to the *world*, they actually talk about how one model relates to another, possibly more complete or accurate, *model of the world*. [14] That is, they are concerned with the relation between the structure $S=<U, O, R>$ and another structure $S'=<U', O', R'>$ which is, after all, a structure (and hence a model) as well. To avoid an infinite regress (and to avoid the somewhat postmodernist consequence that theories always just relate to other theories but never to anything in the world), this regress has to come to a halt. This means one has to believe that at some point structures 'just sit out there in nature'; that is, one has to assume that target systems exhibit structures in roughly the same way it posses simple properties such as colour or location.

Structuralists seems to be quite willing to do so. Though they do not (to my knowledge at least) address this issue explicitly, for the most part they take structures to be unproblematic and think about them in roughly the same way one thinks about simple properties. Systems 'just have' structures and it is the task of scientific research to 'pick them up'.

This assumption, though seemingly plausible, turns out to be mistaken upon closer examination. In what follows I will argue that structures do not just 'sit there' waiting to be picked up. Objects per so do not have a structure at all. Structures have to be ascribed to the system. This, however, involves the construction of a *physical design* (by which I mean roughly something like a model in the straightforward scientific sense). As a consequence, the structuralist notion of a scientific model becomes untenable.

---

[14] This way of putting it is due to Nancy Cartwright and Mauricio Suárez's 'In Favour of Theories as Tools: A reply to French and Ladyman' (unpublished manuscript).

4. Structures Are not for Free

(1) The argument

The argument in this section is a bit intertwined and at times lengthy. For this reason I present a summary of the argument in this subsection, which, I hope, will make it easier to follow the main line of argument later on.

In the previous section I have shown that the stucturalist view must assume that target systems are structured entities; that is, it has to presuppose that they exhibit a physical structure $S_p = \langle U_p, O_p, R_p \rangle$. Proponents of the view tend to see this assumption as wholly unproblematic - so unproblematic in fact that it is usually not even explicitly mentioned. It is the aim of this section to show that this is mistaken. I do not deny that target systems *can* exhibit structures, but I deny that this comes for free. More specifically, my claim is that systems *per se* do not exemplify any particular structure at all. Structures do not just sit 'out there', neatly exposed, waiting for us to portray them. 'And so what?', the proponents of the structuralist view might now counter, 'what is the problem with this?'. The problem, I think, is the following. A system is a 'compact' and unstructured entity and we have to carve it up in order to impose a structure on it. Structures do not really exist until the scientist's mind actually 'creates' them or, to put it in a less pretentious way, ascribes them to a system. More specifically, what we have to do is to identify a set of individuals, which can serve as the domain of the structure and then identify a set of relevant relations and operations on this set. And this may not be a straightforward process. It characteristically involves a host of modelling assumptions (by which I mean, among others, simplifications, idealisations, approximations, the imposition of boundary conditions, and many other 'distorting' processes), and only after having presented the system in the light of such assumptions can we ascribe a structure to it. Structures are not 'ready made' but result from a way of taking, or demarcating, the system. In some cases this can be fairly simple, but normally the assumptions involved are far from trivial. The examples I discuss in the remainder of this section illustrate this point.

A more concise way of putting this point is as follows. What we have to come up with is a *physical design* of the target system. What I mean by physical design is, roughly speaking, a structured version of the target. More specifically, a physical design is an imagined entity consisting of clearly demarcated parts that possess certain properties, stand in certain relations to other parts of the system, and satisfy certain claims and that bears a clearly specifiable relation to (some aspects of) the target system; that is, it is an imagined physical item which is equipped with an exactly describable 'inner constitution' consisting of a web of properties and their interactions. By properties I mean substantive physical properties such as mass, charge, hardness, etc. The salient point in all this is the following. This 'inner constitution' is such that it gives rise to a structure. That is, the parts of the system compose a set U which is the domain of a structure and the relations between the parts are specified in a way that for any n-place relation all the parts that satisfy a relation can be grouped together in n-tuples, and similarly for operations. Given this, U along with the relations and the operations in which its elements enter constitutes a structure.

As a straightforward example take the solar system. What we are dealing with when investigating the system is not the system itself, but a physical design of the system: A composite entity consisting of ten perfect spheres with a homogenous mass distribution; one of these spheres has a privileged status in that it posses almost all the mass in the system, the other nine spheres are orbiting around it, interactions take place only between this big sphere and the other spheres but not among them, and the strength of the interaction is proportional to the inverse square of their distance.

This is still a rough characterisation of physical designs, but it is sufficient to indicate what I have in mind when I talk about physical designs; and the concrete examples I discuss in the following will provide further illustration of what all this amounts to. (I also should note that from this description of physical designs it is obvious that they are closely related - not in concrete detail but in spirit - to Cartwright's prepared descriptions (1983, 133-4), Achinstein's theoretical models (1968, 209), Suppes' physical models (1970, Ch.2 p. 9), and Hesse's analogical models (1963, Ch.2).)

The punch line of this is the following: Physical designs have structures, target systems do not. The solar system *per se* has no particular structure at all. What are its basic entities that constitute the elements of the domain of the structure? The obvious reply is that these are the planets. But this is by no means necessarily so. Why not consider the individual atoms in the system as basic entities? Or why not take a 'Polish' stance and take also the mereological sum of some planets as basic objects? And similarly for the relations between the objects: The choice to neglect all interactions but the ones between the sun and the planets is by no means the only possible choice. For short, there are many possibilities and it is just the physical design that demarcates the phenomena (or carves them up) in way that gives rise to a structure, because just under this design the system consists of clearly defined and identifiable parts and relations between them. For short, no physical design, no structure – this is the first claim that I try to substantiate in the remainder of this section by dint of several concrete examples.

If true, this renders the view that scientific models *are* nothing but structures untenable. Given that structures cannot be tied to reality without the aid of a physical design, why then exclude it from being part of the model? Models are representative devices, so it seems natural to consider everything needed to do perform this function as a part of the model. From this point of view then, there is no reason to exclude physical designs from being part of the model and there is no intelligible reason to restrict the notion of a model to structures. Therefore, a notion of substantial modelling (which is what I try to capture with my physical designs) is not merely a 'friendly amendment' to the structuralist point of view (as van Fraassen 1991, 13-15, urged against Cartwright) but an integral part of any workable conception of a scientific model. Obviously, one can always play a verbal trick and reserve the term for the bits and pieces one particularly likes, but at least in this case this seems to be besides the point.

Proponents of the structuralist view might now argue that all this, though true, does not threaten the core of their account because the whole process of carving up the system is essentially an engineering problem that has not much to do with a *philosophical* analysis of models. I think that this is mistaken for (at least) two reasons: First, it would be a rather strange attitude to ban certain essential processes of scientific research from the realm of philosophical analysis. I take it that it is important to understand what assumptions we make when we model a system; that is, we have to understand what we consider to be basic entities of our domain and what kind of idealisations, approximations, and other simplifications we make. However, I admit that this may be a matter of philosophical taste and I will not build on this argument in what follows.

There is another reason why we cannot omit physical designs from our analysis of representation. It is often possible to carve a target system in different non-equivalent ways. In the case of the solar system the alternative ways may seem a bit contrived, but in many cases there are several equally legitimate ways to delineate a system, which give rise to different non-isomorphic structures. Moreover, the choice of certain modelling assumptions is often dependent on pragmatic considerations as well as on the context of investigation, and as a consequence the structures we are dealing with, though realised in a given system, are dependent on the observer's decisions, purposes, previous knowledge and experience.

This effectively undercuts any attempt do dismiss physical modelling as philosophically irrelevant. If a system possessed just one unique structure one could argue that it is not of philosophical interest to study the processes by which this structure is uncovered. But given that there is no such thing as *the* structure of a system and if a system may exhibit many different, non-isomorphic structures depending on the physical designs we decide to use, we cannot dismiss these as irrelevant. Structures piggy-bag on physical designs and without them they cannot be connected to a target at all. For this reason they are an essential part of a workable account of representation.

In passing, I would like to make a few comments on Newman's theorem. This theorem has repeatedly been used as an argument against structuralism and since my points (seemingly) bear some similarities to it, it is worthwhile to point out what the difference are. In 1928, the Cambridge mathematician M. H. A. Newman proved at theorem stating (roughly) that any set can be structured in any way you like subject to cardinality constraints. In his own words (1928, 144): 'Any collection of things can be organised so as to have the structure W, provided there are the right number of them.' This sounds very much like my claim that systems may exhibit different, non-isomorphic structures. However, though the two arguments pull in the same direction in the sense that they aim at undermining a structural essentialism holding that every systems has exactly one characteristic structure, there are important differences between the two.

My argument is, at once, stronger and weaker than Newman's. The proof of Newman's theorem turns on the fact that relations are understood extensionally in set theory (i.e. an n-place relation is taken to be no more than a set of ordered n-tuples). Hence, given a bunch of objects, we can structure them in any way we like just by putting the objects into ordered n-tuples; there are no constraints on this procedure other than that we need enough of them. (For instance, from this point of view it may be perfectly meaningful to say that the city of London has the structure of a Semi-Riemannian manifold.) No physical constraints are taken into account; this 'grouping together' of objects (which creates the relations and hence the structure) is a purely formal procedure that pays no attention the nature of the objects involved, and the relations so created do not need to have any physical reality (in the sense that they have any influence, observable or not, on the system's behaviour).

It is at this point that the two arguments diverge. What I want to argue is that a system can exhibit different *physically relevant* structures, i.e. structures that reflect the essence of a phenomenon and are not merely formal constructs. (I am aware of the fact that this is a somewhat vague characterisation, but I think there is no general description of what it means to capture the essence of a phenomenon, so I rely of the subsequent examples to clarify what I have in mind.) In this sense, my claim is stronger than Newman's. But at the same time it is weaker since this restriction to physically relevant structures drastically narrows down the freedom of choice. Among the great many structures compatible with cardinality constraints only a few will be physically relevant. Realistically, the scientist has to choose between a certain (rather limited) number of ways to carve up a system. Though there are some reasonable options, there normally is not a huge number of them – not anything goes. Which one of these we choose to use is often dependent on the context, the purpose of the investigation and other pragmatic factors. However, to discuss how these pragmatic factors determine certain choices is beyond the scope of this paper. The point I want to argue for is that these alternative structures do exist; it is another project to discuss what factors determine their choice and how they do so.[15]

Summing up, my two main claims are: (1) We have to choose a physical design before we can meaningfully say that a system has structure, and (2) there are several equally good physical designs for a particular target system and for this reason there is no such thing as *the* structure of a system. In the remainder of this section I will substantiate these two claims by offering several concrete examples.

(2) A starting point: objects of everyday experience

To begin with, consider physical objects of our everyday experience. They are not presented to us in 'analysed' form; they do not normally consist of neatly defined parts. In most cases we face a 'compact entity' and we have to 'cut it up', to carve it, before structures can be discussed. And this may not be an easy task. To get the flavour, consider a hum-drum example: What is the structure of the Eiffel Tower? We could, for instance, take all the rivets that keep the construction together as the objects and their spatial arrangement as relations. Or we could take the intersections of the iron bars as objects and the forces between them as relations. Or to make it simpler, we could take the three levels of the tower as individuals and consider them as structured by the relation 'higher than'. Or we could divide the tower into a North and a South part.

This small list by no means exhausts the possibilities, there are many more ways to carve the tower, and each may have its legitimacy. The first two may be of interest to construction engineers, the third to the provider of elevators, and the fourth to those who care about rust protection. The conclusion is: The object only has a structure given a certain physical design.

[15] Eric Peterson in his PhD thesis (UCSD, forthcoming) provides an interesting case-study form economics and finance in which he shows how pragmatic factors influence the choice of structure.

There is no such thing as *the* structure of the Eiffel Tower. We have come up with a design first; only then structures enter the scene.

(3) A continuation: the methane molecule

The Methane molecule (CH_4) consists of one Carbon and four Hydrogen atoms. Since Carbon and Hydrogen have the same electro-negativity the four Hydrogen atoms form a regular tetrahedron at the centre of which we find the Carbon atom. Hence, the space occupied by the molecule is a regular tetrahedron (which is to say that the compound has a tetrahedrical shape). In many scientific contexts (e.g. collisions or the behaviour of a molecule *vis a vis* a semipermeable membrane) only the shape of the molecule is relevant, that is, we can forget about the carbon atom sitting at the centre of the molecule. This naturally raises the question: What is the structure of (the shape of) the Methane molecule?

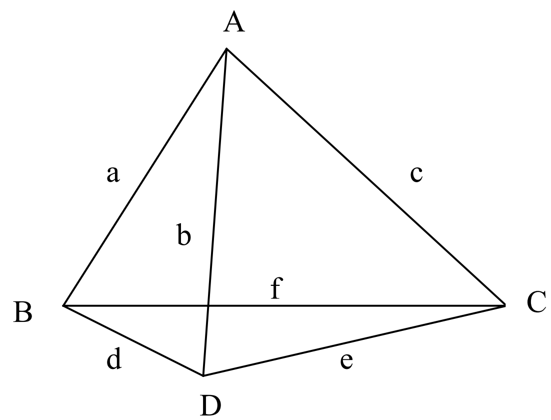

Fig 2. Tetrahedron.

To apply our notion of structure we need a set of basic objects and relations on it (leaving aside operations for the time being). Now the difficulties begin (see Rickart 1995, 23, 45). What are the objects that constitute the domain of the structure? A natural choice seems to take the corners (vortices) as the objects and the lines that connect the vortices (the 'sides' of the tetrahedron) as the relations holding between the objects (the relation then is something like 'being connected by a line, or more formally Lxy = 'x is connected to y by a line'). We end up with the structure T_V which consists of a four-object domain, the four vortices A, B, C, and D, and the relation L which has the extension {(A, B), (A,C), (A,D), (B,C), (B,D), (C,D)}.

However, this is neither the only possible nor the only natural choice. Why not consider the lines as the objects and the vortices as the relations? There is nothing in the nature of vortices that makes them more 'object-like' than lines. Following this idea we find the structure T_S with a domain consisting of the six sides a, b, c, d, e, and f on which the relation I (Ixy = 'x and y intersect') is defined. It has the following extension {(a,b), (a,c), (a,d) (a,f), (b,d), (b,e), (c,e), (c, f)}.

Hence, we need a modelling assumption specifying what the objects and the relations are before we can tell what the structure of the (shape of the) molecule is; that is, we need a physical design of the molecule before structures enter the scene. The tetrahedron *per se* has no structure at all; it just has a structure with respect to a certain description, namely one that specifies that the vortices are the individuals in the domain of the structure and the lines the relations, or vice versa.

But which one of these structures is *the* structure of the (shape of the) molecule? I think there is no fact of the matter. There is no reason why one structure should be privileged - both are based on 'natural' features of the tetrahedron. Hence, our molecule has (at least) two structures. And the important point is yet to come: These structures are not isomorphic. This follows immediately from the fact that the domains of T_S and T_V do not have the same cardinality.

The upshot of this example is that even a simple geometrical object like the methane molecule does not posses a structure *per se* and we can, using different designs, attribute different *non*-isomorphic but equally legitimate structures to it.

This is not the end of the story yet. We can now replace one of the Hydrogen atoms by another Halogen, Chloride say, and obtain CH_3Cl. Since Chloride does not have the same electro-negativity as Hydrogen the new compound has the shape of an *irregular* tetrahedron, i.e. one whose sides do not have the same length. It is obvious that this tetrahedron still exemplifies T_S and T_V. However, due to its irregularity, it exemplifies further structures: Take again the vortices to be the objects and define the relations between two vortices to be the length of the side connecting them. We then get six different relations (assuming that all sides are of different length) whose extension comprises just one pair each. This new structure, call it T_L, is obviously non-isomorphic to both T_V and T_S. Hence we have ascribed three equally good structures to the irregular tetrahedrical shape of CH_3Cl just by adopting different modelling assumptions, and I don't doubt that with some ingenuity one can find many more.

Finally, it goes without saying, the same is true for more complex objects like cubes or other polyhedrons, or more generally for any object consisting of lines (not even necessarily straight) that intersect at certain points.[16]

What we learn from this example, I take it, is that molecules do not exhibit one particular structure in any obvious way. Objects of this sort may be analysed in more that one way with respect to structure. And this is by no means a peculiarity of the above example. The argument only relies on very general geometrical features of the shape of a molecule and these can easily be carried over to any kind of object.

[16] For this reason Rickart (1995, 19) distinguishes between systems and structures. A system, on his definition is 'any collection of interrelated objects along with all of the *potential* structures that might be identified with it [...] a system is only "potentially" structured. It will exhibit a structure as soon as any of the potential structures is made explicit.'

(4) Logistic growth in ecological models

Suppose we are interested in the growth of a population of some particular species, wasps say. One of the earliest, and by now famous, models has it that the growth of the population is given by the so-called logistic map,

$$x' = Rx(1-x),$$

where x is the population density in one generation and x' in the next, R is the growth rate. For the sake of the argument assume that this equation describes the situation correctly. From a structuralist point of view this amounts to saying that the structure that is defined by the logistic map, S_L, is isomorphic to the structure S_W of the system (i.e. the population of wasps) and that S_L represents the population of wasps.

A closer look at how the model really works quickly reveals how far off this is. As Hofbauer und Sigmund (1998, 3) point out, in many ecosystems thousands of different species interact in complex patterns and even the interactions between two species can be quite complicated, involving the effects of seasonal variations, age structure, spatial distribution and the like. But nothing of this is visible in the model. No interaction with any part of the ecosystem is explicitly built into the model. It is just the net effect of all interactions that is accounted for in the last term of the equation ($-Rx^2$), which reflects the fact that a population cannot grow infinitely due to restrictions imposed by the environment. Hence, all actual interactions are 'idealised away'.

Once this is done, one has to define the objects of the structure. An obvious choice would be individual wasps. But one readily realises that this would lead to huge and intractable sets of equations. The 'smart' choice therefore is to take generations rather then individual insects as objects. But this is not enough yet. We have to assume furthermore that the generations are non-overlapping, reproduce at a constant average rate (reflected in the magnitude of R) and in equidistant discrete time steps.

The thrust of my argument is clear by now: There is not structure just 'sitting there' waiting to be picked up. The analysis of certain structures into the target system may be a quite difficult and laborious task and a given system may be analysed in more than one way, depending upon which aspects of the system one wishes to emphasise. The system does not have any particular structure at all unless a certain physical design is attached to it. Before we do not decide to take generations as basic units of study, make time discrete, neglect the interactions with other parts of the ecosystem, and so on, one cannot meaningfully say that the target system exhibits a structure.

(5) Summing up

The main conclusion to be drawn form this is that systems *per se* do not exemplify any particular structure at all; structures have to be ascribed to a system. Systems just exhibit a structure given physical design; or to put it differently, if they are modelled in a certain way: No modelling assumptions, no structures! These assumptions can be rather modest as in the case of the Methane molecule where they basically consist in specifying what the basic objects are, or they can be substantial as in the population dynamics case, but in either case we cannot do without. We have to present a design of the target before structures can even enter the scene. In this sense structures are the end product of an investigative endeavour and not the starting point.

5. Conclusion and Outlook

(1) Representation as a two tiered concept

What has been achieved? The argument developed in this paper is kind of *reductio ad absurdum* of the structuralist conception of models. I started with the observation that models do represent things in the world and then asked whether the stucturalist conception of models can account for this fact. My conclusion is that it can't. Structures *per se* do not represent. So we have to add some 'stuff' in order to endow structures with representative power. But what is this stuff? A rather lengthy discussion has led me to the conclusion that the 'stuff' we have to provide is a physical design. It is this design to which the theoretical structure (the structuralist's model) is attached via a process of matching (which is very similar to reference-fixing in the case of language). Given that this design is essential to achieve representation, I cannot see why it should be excluded from the entity we call model. Models are representative devices and it is natural to define them in such a way that they comprise everything that is need to achieve this goal. Therefore structures by themselves (even augmented with an isomorphism claim) are not models. A models is, roughly speaking, a structure plus a physical design plus a matching process.

But this is hardly structuralism any more. Though I do not deny that structures are part of the picture, they are just a part among others.[17] The physical design is just as important as the structure itself. Models are complex entities and cannot be reduced to mere structures. And for this reasons physical designs are not merely a friendly amendment to the structuralist view. Schematically, the new picture I suggest looks as follows:

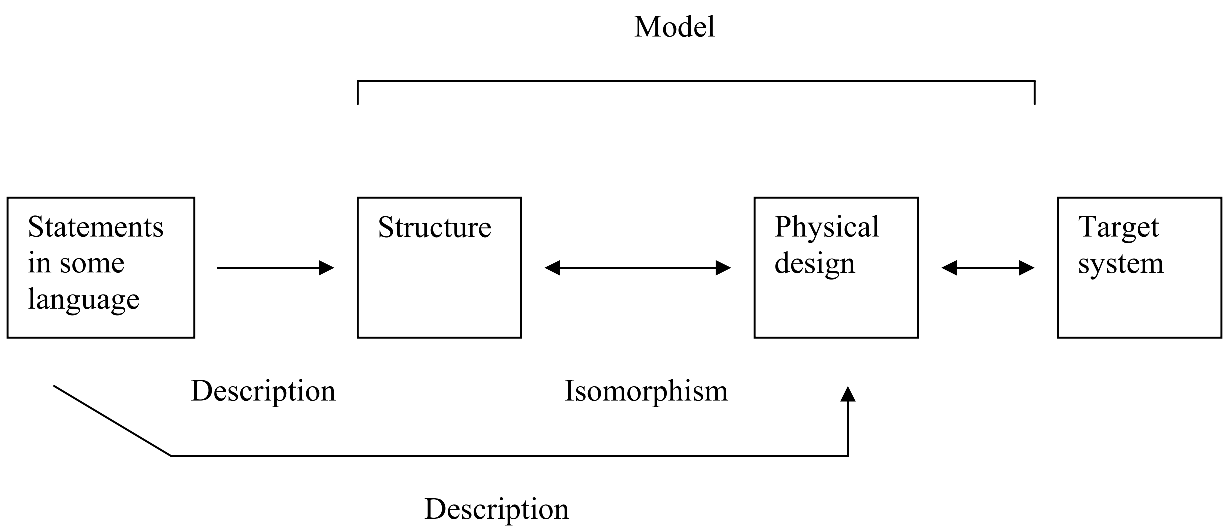

Fig. 3: A new schema for representation.

This raises two questions (at least): First, where does representation take place? Is it between the structure and the physical design or is it between the physical design and the target? Intuitively we think that representation consists of one thing standing for another, but in the above diagram we have a chain of things relating to each other in different ways and it is not clear where to locate the representation relation. Second, some may have noticed that the arrow in

[17] This describes the case of mathematised models (which was the issue at stake in this paper). In the case of non-mathematised models, however, the structure is missing and the physical design by itself constitutes the model; scale models (e.g. of planes or bridges) are a point in case.

the above figure that connects the physical design to the target system is not labelled. This is not merely a sin of omission. In my characterisation of a physical design (in Sec. 4.1) I have mentioned somewhat miraculously that it 'bears a clearly specifiable relation to (some aspects of) the target system'. What does that mean?

The answers to these questions are everything but obvious and in the following I shall just briefly indicate what seems to me to be the right way to look for solutions. As to the second question, I take it that there is some form of representation involved; in some sense or another the physical design represents its target system. But what precisely is the relation a physical design bears to its target? This is an important question that deserves more attention. How does one object (the design), represent another object (the target). What notion (or notions) of representation is involved? While word-to-object representation has received a tremendous amount of attention in analytical philosophy, object-to-object representation has, after some pioneering work by philosophers like Achinstein (1965, 1968), Black (1960) or Hesse (1953, 1963), fallen into oblivion. I suggest that this should be changed. Object-to-object representation is extremely important and it should be the subject of intensive study in the future. What Goodman wrote in 1968 can still be a motto today, at least as far as the philosophy of science is concerned: 'Systematic inquiry into the varieties and functions of symbols has seldom been undertaken. Expanding investigation in structural linguistics in recent years needs to be supplemented and integrated with intensive examination of *nonverbal symbol systems* […] if we are to achieve any comprehensive grasp of *the modes and means of reference* and of their varied and pervasive use in the operations of *understanding*' (Goodman 1968, xi, all emphases mine).

This suggests a (partial) answer to the first question. There does seem to be one single representational relationship. Models are complex entities and there are different relationships between different parts of the model and the world that are all representational in some sense.

The above figure proposes a two tiered concept of representation: A structure represents a physical design due to its being connected to it by a process similar to reference-fixing in the case of language and due to isomorphism; and a physical design represents its target system in way that still needs elucidation (but which can be expected to be related to methods used in iconic or analogical modelling). In either case representation takes place but it is of a wholly different kind. What the nature of these relations is remains to be discussed, but one thing we know for sure by now: We are barking up the wrong tree if we are looking for *the one and only* relation of representation.

Finally I should emphasise again that my strategy has been to undermine the structuralist view 'from inside'. The structuralist conception of model has come under attack from different sides as early as 1968 when Achinstein (Ch. 8, esp. p. 249) doubted their fruitfulness for scientific practice (he referred to them as 'mathematical models') and pervasive criticisms along the same line have been put forward in the 'models as mediators' project. Though I agree on this, I have not used *these* arguments. It is *not* my point that structuralism cannot account for scientific practice or for certain historical episodes; my line of criticism is that its models have no representative power and by trying to endow them with it much more than structuralism can provide is needed. For short, the dilemma is the following: Either one is a structuralist and is prepared to accept that models do not represent – i.e. that one is doing 'physics without physics', as Jeff Ketland (2002) puts it – or one has to bring physical designs back into the picture and thereby abandon stucturalism proper, because this is assuming much more than even a modest structuralist can account for.

(2) The aim and purpose of structuralism

What has gone wrong? Why has the structuralist conception of models been able to gain so much ground if it is not able to account for one of the most basic features of scientific models, namely representation? The answer to this question is, paradoxical as it may sound, that it has not been designed to do so in the first place. The structuralist view prima facie construes models *as models of theory*, and not as models of anything in the world (see e.g. Suppes 1970 pp. 1-6 and 1-9). As Suppes (1988, 254) remarks, to study formal models is the best method to gain insight into the structure of a complex theory since its syntax is often intertwined and offers little insight into the nature of the theory. I entirely agree with that. But to gain insight in the nature of a formal theory and to device a representative tool are two different issues. From this point of view the failure of structuralism in the field of representation is little surprise. Why claim in the first place that structuralism is an account of scientific representation? Why not use it as a tool for formal analysis and admit that representation is a different issue that needs to be tackled with different methods? The structuralist view of representation is nothing but a stopgap. Once this is realised, the way is free for a more constructive debate of the topic.[18]

Acknowledgements

[18] One of the reasons why this happens may be mere habit. One gets so used to the mathematics that one suddenly forgets that there are other elements to a complete model. Here is a telling example. In the preface to the second edition of their book on evolutionary games Hofbauer and Sigmund write: 'In our former book, it took us 150 pages of biological motivation to tentatively introduce the notion of a replicator equation. This is no longer warranted today: replicator dynamics is a firmly established subject … and our old volume definitely looks dated today.' (1998, xi). This boils down to: Now we can do maths without bothering about the biology too much any more – and that is how everything that is not structural gets put out of sight.

I would like to thank Nancy Cartwright, Matthias Frisch, Carl Hoefer, Jeff Ketland, Julian Reiss, and Mauricio Suárez for comments on earlier drafts and/or for helpful discussions; thanks to Valeria Mosini for an interesting discussion about molecules. An earlier draft of this paper has been presented on 13May 2002 at the Expert Workshop “Measurement in Economics and Natural Sciences” in the Tinbergen institute in Amsterdam and I am grateful for comments and suggestions made by participants in the debate.

References

Achinstein, Peter (1965): Theoretical Models. In: British Journal for the Philosophy of Science 16, 102-20.

- (1968): Concepts of Science. A Philosophical Analysis. Baltimore.

Bell, John and Moshé Machover (1977): A course in Mathematical Logic. Amsterdam a.o.

Black, Max (1960): Models and Archetypes. In: Black, Max: Models and Metaphors. Studies in Language and Philosophy. Ithaca / New York, 219-43.

– (1970): How do Pictures Represent? In: Ernst H. Gombrich, Julian Hochberg and Max Black: Art, Perception, and Reality. Baltimore and London 1970, 95-130.

Boolos, George S. and Richard C. Jeffrey (1989): Computability and Logic. 3rd ed., Cambridge.

Bueno, Otávio (1997): Empirical Adequacy: A Partial Structures Approach. Studies in History and Philosophy of Science 28, 585-610.

- (1999): What is Structural Empiricism? Scientific Change in an Empiricist Setting. Erkenntnis 50, 59-85.

Cartwright, Nancy (1983): How the Laws of Physics Lie. Oxford.

- (1999): The Dappled World. A Study of the Boundaries of Science. Cambridge.

Costa, Newton da and Steven French (1990): The Model-Theoretic Approach in the Philosophy of Science. Philosophy of Science 57, 248-65.

- (2000): Models, Theories, and Structures: Thirty Years On. Philosophy of Science 67 (2000) Proceedings, S116-127.

Fraassen, Bas C. van (1980): The Scientific Image. Oxford.

- (1989): Laws and Symmetry. Oxford.

- (1991): Quantum Mechanics. Oxford.

French, Steven (2000): The reasonable effectiveness of Mathematics: Partial Structures and the Application of Group Theory to Physics. Synthese 125, 103-20.

French, Steven and James Ladyman (1999): Reinflating the Sematic Approach. International Studies in the Philosophy of Science 13, 103-121.

Giere, Ronald N. (1988): Explaining Science. Chicago.

Goodman, Nelson (1976): Languages of Art. 2nd ed. Indianapolis and Cambridge.

Hesse, Mary (1953): Models in Physics. In: British Journal for the Philosophy of Science 4, 198-214.

- 1963: Models and Analogies in Science. London / New York.

Hodges, Wilfrid (1997): A Shorter Model Theory. Cambridge.

Hofbauer, Joseph and Sigmund, Karl (1998): Evolutionary Games and Population Dynamics. Cambridge.

Ketland, Jeff (2002): Physics Without Physics: The Newman problem for Structuralism. Unpublished Manuscript.

Kroes, Peter (1989): Structural Analogies between Physical Systems. British Journal for the Philosophy of Science 40, 145-54.

Machover, Moshé (1996) Set Theory, Logic and Their Limitations. Cambridge.

Newman, M. H. A. (1928): Mr Russell's 'Causal Theory of Perception'. In: Mind 37, 137-48.

Putnam, Hilary (1981): Reason, Truth, and History. Cambridge.

Redhead. Michael (2001): The Intelligibility of the Universe. In: Anthony O'Hear (ed.): Philosophy of the New Millennium. Supplement to 'Philosophy', Royal Institute of Philosophy Supplement 48, 73-90.

Rickart, Charles E. (1995): Structuralism and Structure. A Mathematical Perspective. Singapore a.o.

Shapiro, Stewart (2000): Thinking About Mathematics. Oxford.

Steiner, Mark (1998): The Applicability of Mathematics as a Philosophical Problem. Cambridge MA and London.

Suárez, Mauricio (1999): Theories, Models, and Representations. In: L. Magnani, N. J.Nersessian, P. Thagart (eds.): Model-Based Reasoning in Scientific Discovery. New York, 75-83.

Suppe, Frederick (1989): The Semantic Conception of Theories and Scientific Realism. Chicago.

Suppes, Patrick (1960a): A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences. In: Suppes, Patrick: Studies in the Methodology and Foundations of Science. Selected Papers from 1951 to 1969. Dordrecht 1969, 10-23.

- (1960b): Models of Data. In: Suppes, Patrick: Studies in the Methodology and Foundations of Science. Selected Papers from 1951 to 1969. Dordrecht 1969, 24-35.

- (1970): Structures in Science. Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford.

- (1988): Representation Theory and the Analysis of Structure. Philosophia Naturalis 25 (1988), 254-68.