2 Models and Modelling

The term ‘model’ is used in many different situations and in many different ways. This causes some difficulties when trying to discuss the uses of models – what Wartofsky calls the “Model Muddle”:

> “*The symptom of the (model) muddle is the proliferation of strange and unrelated entities which come to be called models.” [462] p.1*

He goes on to identify the root cause of this proliferation, the fact that “*anything can be a model of anything*” *[462]* p.4 provided only that:

> “*… every case of being taken as a model involves a restriction with respect to relevant properties.” [462] p.6.*

Wartofsky does not mean that anything can be used as a model for anything for any purpose (for example, a robin used as a model of urban development), but that one can not rule out *A* being used as a model for *B* just by the nature of *A* and *B* alone. There may be some properties of *A* and *B* which when selected as relevant (for some purpose) will legitimise *A* being taken as a model of *B*. This is implicit in his statement that

> “*The limits on modelling are therefore the limits of our conceptions of what respects are relevant to what purposes.” [462] p.6*

Of course, these conceptual limitations might be quite strong! We might be limited in our conceptions by our experience of the world and our reasoning about it, for example. However, our conceptions of the world and hence which things are relevant to which might be very strange – it is conceivable that for any objects or systems, *A* and *B*, *someone* might be have a (possibly mistaken) conception of the relevance of their properties so that that person might use *A* as a model of *B* (despite the fact that the rest of us might see this as pointless).

Thus modelling can be seen as triadic relation, $M(S, A, B)$ – *S* takes *A* as a model for *B*, by finding or choosing relevant properties of both, $R(A)$ and $R(B)$ such that $R(A) \subseteq R(B)$. This is illustrated in figure 1. To avoid repetition I will refer to the object or system being modelled as simply the *object*, unless this is unclear.

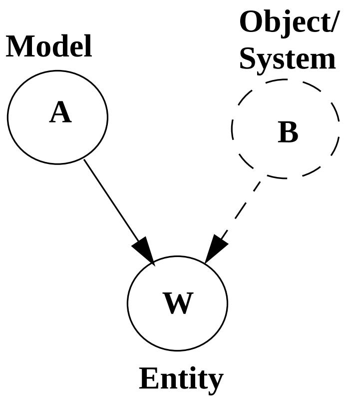

Figure 1. Entity, W, using A as a model of B

Apostel relativises this relation further, by making explicit the purpose (or 'modelling goal') of the person in using a model. Thus he writes the basic relation as $M(S, P, A, B)$, meaning that "*The subject S takes, in view of the purpose P, the entity A as a model for the prototype B.*" [16] p. 4.

These are, however, very general definitions which do not tell us much about the nature of models as actually used by people; how to distinguish between useful and less useful models, or how they should use them. To do this one has to narrow the context down and subsume some of the necessary or accepted limitations on our conceptions of relevance in the places where models are used in a deliberate manner. The area where this has been most studied is in the practice and theory of science.

Models in the philosophy of science are usually distinguished from theories. However, the distinction is not always clear cut and some authors slip from one to the other. Three examples of philosophers who use the terms differently are: Apostel [16] who allows a theory to be a model, Hesse [216] who argues that a theory is meaningless if it does not *include* a model and van Frassen who says "*To present a theory is to specify a family of structures, its models…*" [161] p. 64). On the whole the use of the term 'model' implies something relatively concrete, possibly visualisable, approximate, specific and tractable whereas a 'theory' implies a more general, abstract and reliable construction. In the social sciences it is more common to use the term 'model' for things that would be called theories in the physical sciences (according to Braithwaite, 1962, who deplores the practice).

The distinction between models and theories is clearest in those who take a ‘model-theoretic’ approach. Here a strong analogy is made with the Tarskian logical semantics – the theory is likened to a logical theory describable in terms of a set of axioms along with an inference procedure and a scientific model is likened to a logical model, that is a formal structure which satisfies the theory. According to DaCosta and French in [129] this approach was first introduced by Beth [61, 62, 63] and Suppes [430, 432]. Others wish to deliberately conflate these entities in the search for a more basic structure, from which the ‘headline’ entities might be built (e.g. Giere [176] and Wartofsky [462]). My approach will tend towards the later, since I am not concerned here with the truth of such entities but with the complexity that arises as a result of the process of modelling. However, I will maintain a clean distinction between the syntactic and semantic parts of the modelling apparatus.

2.1 Some Types of Models

I will briefly outline some of the types of models that have been identified in the philosophy of science literature, before going on to discuss some of the relations between them in the next subsection (Section 2.2).

Firstly, models may be categorised according to the medium in which they are expressed. In this way we have physical models including scale models (the first type of model of Achinstein in [6] and the second sense of model of Suppe in [427]) and Irving Fisher’s model of the money supply using coloured liquids (as documented by Morgan in [328]); mathematical models[10] where the language of mathematics is the primary medium for defining the content of the model; computational models where the model is encoded as a computer program with the computation somehow modeling the object process; and linguistic models where normal (or pseudo-formal) spoken or written language is used to specify the model[11].

Secondly models are distinguished according to whether the term model applies to the sentences and formal theories that describe the model or to the referent of those

10.Unfortunately ‘mathematical model’ has two connotations: that of being expressed in terms of mathematics and that of being formal in contrast to analogical models.

11.There are also, of course, logical models but, in order to prevent confusion I will only use this phrase to denote the sort of models found in metalogic.

sentences. The former have been called ‘syntactic models’ or (more precisely) the syntactic approach to models (e.g. Ayer [32] and Duhem [143]). The later ‘semantic models’ or the semantic approach to models is where the models are the entities described by (or that satisfy) the description – these can be further divided into those which envisage these semantics to be akin to the semantics in metalogic where a formal structure is a model when it satisfies the model specification (e.g. Sneed [417], Suppe [427][12] and Suppes [433]), and those which include a wider, linguistic semantics (e.g. Giere [175]).

In a different way analogical and mathematical models are distinguished. An analogical (or iconic) model involves a (possibly incomplete) analogy between some properties of the model and the object or system being modelled or between the relationships between the properties of the model and the relationships between the properties of what is modelled (as Sellars prefers to define it in [406]). These include the iconic models of Hesse [216] and Suppe [427] as well as model_2 of Cushing in [126]). Morgan [328] further distinguishes between ‘ready made’ and ‘designed’ analogical models. An important aspect of the iconic models is claimed to be that they allow the theory to become predictive, thus Hesse says:

> *“(an iconic model is) any system whether buildable, picturable, imaginable or none of these which has the characteristic of making a theory predictive.”* [216] *p.19*

A mathematical model implies a more formal structure, which may be an approximation or simplification of the underlying theory (including Cushing’s model_1 [126], the ‘theoretical model’ of Achinstein in [5] and Suppe’s first sense of model in [429]). Achinstein’s theoretical model is a type of mathematical model but implies a slightly greater existential commitment, it is somehow an exemplar or expression of a theory. Some authors allow whole series of model types differing in the ascribed existential commitment (e.g. Wartofsky [462]). Kuhn allows models a similar range of epistemic commitment:

> *“Models, … are what provide, the group with preferred analogies or, when deeply held, with an ontology.” [274] p.463*

12.Suppe confusingly calls these ‘mathematical models’.

Cartwright [84] distinguishes between ‘explanatory models’ which have a broad scope and are used to explain the cause(s) of events (but, she claims, are not descriptively true) and the phenomenological laws which are accurate representations of the phenomena but which do not provide satisfactory explanations. In this conception there is a hierarchy of models from the dirty-but-true up to the general-but-inapplicable:

> *“To explain a phenomenon is to find a model that fits it into the basic framework of the theory and that allows us to derive analogues for the messy and complicated phenomenological laws which are true of it.” [84] p.152*

In many of these conceptions (especially those of the semantic approach) a model resides within the framework of a theory. Redhead [372] distinguishes two ways in which this can occur: by simplifying the theory to become more tractable (*impoverished*) or by filling in missing detail to *enrich* the theory.

There are what are called ‘Tinker Toy’ or ‘Floating’ Models (Post, 1975 as reported by Redhead in [372] and Cushing’s model_3 [126]). These are models which are neither verified against known data nor related to a deeper theoretical framework – they just ‘float’ somewhere in between. Despite these having been denigrated in terms of their uncertain status by Post, they have been recently defended by Redhead as a valuable tool to aid the testing of theories [372].

Finally there are: Suppes’ ‘data model’ [431], Kemeny’s “intended factually true description” [255] and the abstract models not intended to be true of an natural thing (Achinstein’s third type of model in [6]). In addition to the above, models could also be divided into the subject areas they originate from or are used in.

Even the above list is by no means comprehensive. Every time a philosopher characterises or defines a model in a different way one could give it a new label and add it to the confusion (Chao in [103] lists no fewer than 30 characterisations). What turns out to be more important than definitions is the different functions and uses that such ‘models’ are put to and the ways in which they are deployed (both singly and in combination) by practising scientists. As Apostel summarises it:

> *“(a model is) … any subject using a system A that is neither directly nor indirectly interacting with a system B, to obtain information about the system B, is using A as a model for B.”* [16]

2.2 Combinations of Models

The knowledge gained from applying analogical models is far from certain, unless you can be sure they are being applied to a situation that is essentially identical to a previous application. The logical empiricists at the beginning of the 20th century wished to put science onto a surer footing: analogies were denigrated as mere heuristics and the focus shifted towards the use of formal systems. As Braithwaite puts it:

> *"Analogy can provide no more suggestions of how the theory might be extended; and the history of science tells us that while some analogical suggestions have led to valuable extensions of the theory, others have led to dead ends" [73] p.230*

The ideal form of scientific knowledge was a formal system of axioms whose logical consequences were to be related to the observed phenomena as directly as possible (figure 2). Thus one had a fairly simple picture with the formal theory being manipulated by syntactic manipulations in order to make predictions about phenomena. As a more descriptive approach to the workings of science has been adopted by many philosophers, there has been a progressive elaboration of this picture.

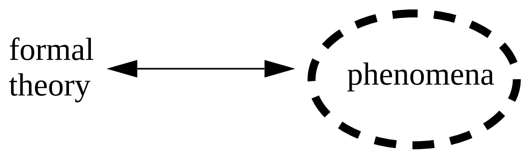

Figure 2. An illustration of the syntactic view of models

Tarskian Logical Semantics provided such an step. Scientific models can be identified with logical models and scientific theories with logical theories. The model, according to this account, is a structure which satisfies the statements that define the theory. This model is then related to the phenomena under study. Two corollaries of this are that different descriptions may identify the same model and that a theory may have several models. This picture of models as semantic entities is illustrated in figure 3 and summarised in the following quote from van Frassen:

> *"… the language used to express the theory is neither basic nor unique; the same class of structures could well be described in radically different*

ways, each with its own limitations. The models occupy centre stage.” [161] *p.44*

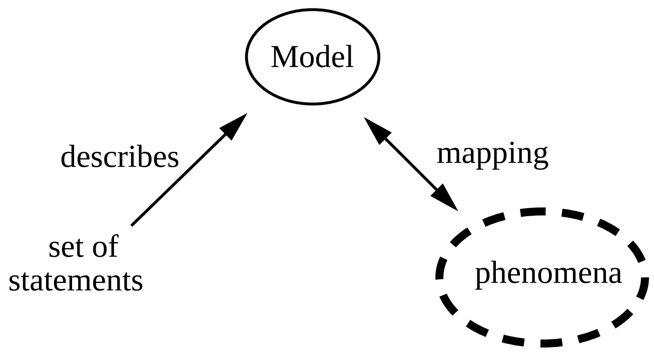

Figure 3. A semantic picture of modelling (from Giere in [176])

Cartwright’s distinction between explanatory and phenomenological laws [84] can be seen as adding a further layer to this picture. The phenomenological laws relate more directly to the phenomena – they have the function of organising the phenomena but do not provide suitable material to construct explanations. Explanatory laws, on the other hand, provide explanations but are not descriptive of the phenomena – they are not literally *true*. This account is interpreted in terms of layers of models by Hughes in [236] and illustrated below in figure 4. In this way the explanatory laws acts as models of the phenomenological laws. Cartwright says:

> “*To explain a phenomenon is to find a model that fits it into the basic framework of the theory and that allows us to derive analogues for the messy and complicated phenomenological laws which are true of it.” [84] p.152*

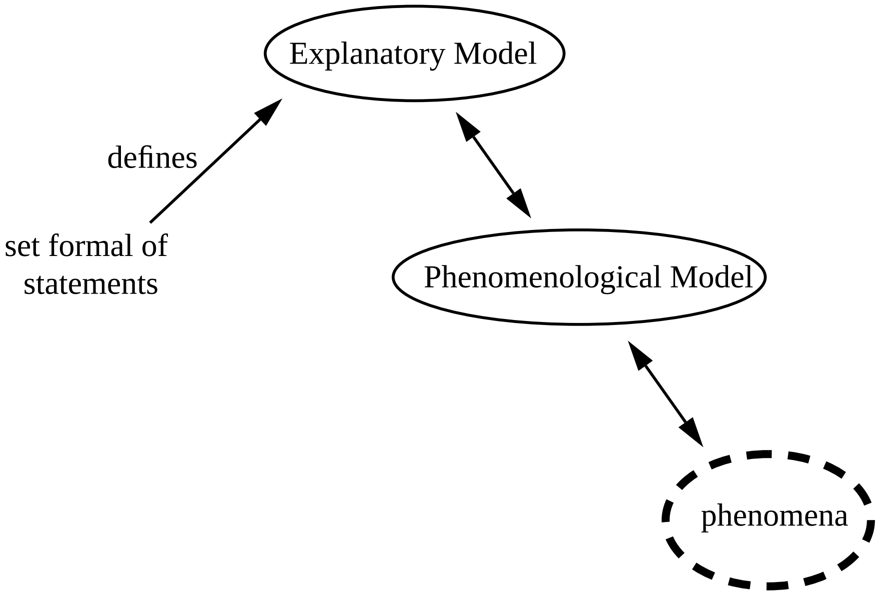

Figure 4. The picture with explanatory and phenomenological models

A further layer is added if one accepts the importance of the distinction between data and phenomena as discussed by Bogen and Woodward in [68]. They point out that data is used more for prediction and the phenomena for explanation. In a real sense we use the data as a 'data model' of the phenomena using the measuring process, as Suppes recognised:

> *"It may be noted that in using an N-tuple as a realisation of the data … we have taken yet another step of abstraction and simplification away from the bewilderingly complex complete experimental phenomena" [431] p.256*

The extended picture is shown in figure 5.

The reality of this distinction make it possible that observations may be theory-laden, since if our data was completely determined by the natural phenomena then the theory-ladenness would not effect our measurement.

Kuhn wrote:

> *"Far more clearly that the immediate experience from which they in part derive, operations and measurements are paradigm-determined. Science does not deal in all possible laboratory manipulations. Instead, it selects those relevant to the juxtaposition of a paradigm with the immediate experience that that paradigm has partially determined." [274]*

Cartwright emphasises rather the process of active adjustment that occurs during the preparation of data:

> *"...when we present a mode of a phenomenon, we prepare a description of the phenomenon in just the right way to make a law apply to it." [84] p.157*

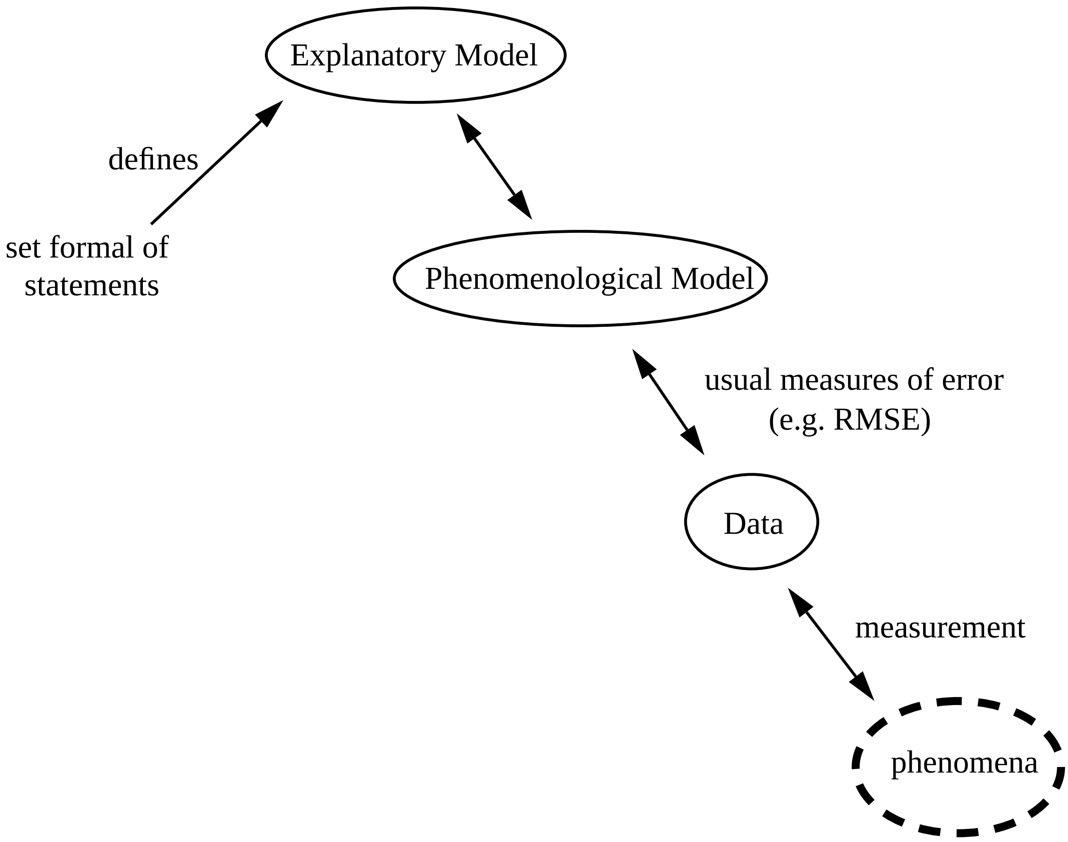

Figure 5. The picture with the data model added

I have built up a picture where a whole series of models come between the phenomena under study and the theory that is supposed to cover them. The relationships between these models are not always clear, especially that between the explanatory and phenomenological models. Cartwright argues that a 'covering-law' account is insufficient, as it frequently does not seem to be the case that one can derive aspects of the phenomena by gradually adding auxiliary hypotheses and facts as steps down this ladder [84]. Thus one comes to a picture of models as objects that somehow *mediate* between the phenomena and the theory - models as mediators [329].

For example, in chemistry the gap between theory and experimental phenomena is being increasingly bridged via the extensive use of computer-based simulations of many interacting atoms [201]. A model of chemical interaction or structure is posited using approximation, simplifications, abstractions, etc. of the underlying (and previously well established) physical and chemical models of pairwise atomic interaction. This model is then encapsulated in a computational model involving many such interacting atoms. The simulations are run to explore the envelope of possible outcomes which are then approximated by mathematical functions. Thus the simulation substitutes for the intractable symbolic inference that otherwise would be necessary. Finally this post-hoc predictive model is compared against the data (or data model) derived from a measurement process applied to the chemical phenomena under study. This is illustrated in figure 6 (which is an adaption of Figure 17 on page 1021 of [201]). I quote

> *"… the core of any model is the approximations, assumptions and simplifications that are used in its formulation. Only an understanding of the basics of a particular model may lead to sensible application or improvement of its performance. … Due to the complexity of the systems of chemical interest, theoretical methods only became of practical interest with the advent and development of computers." [201] p.1021*

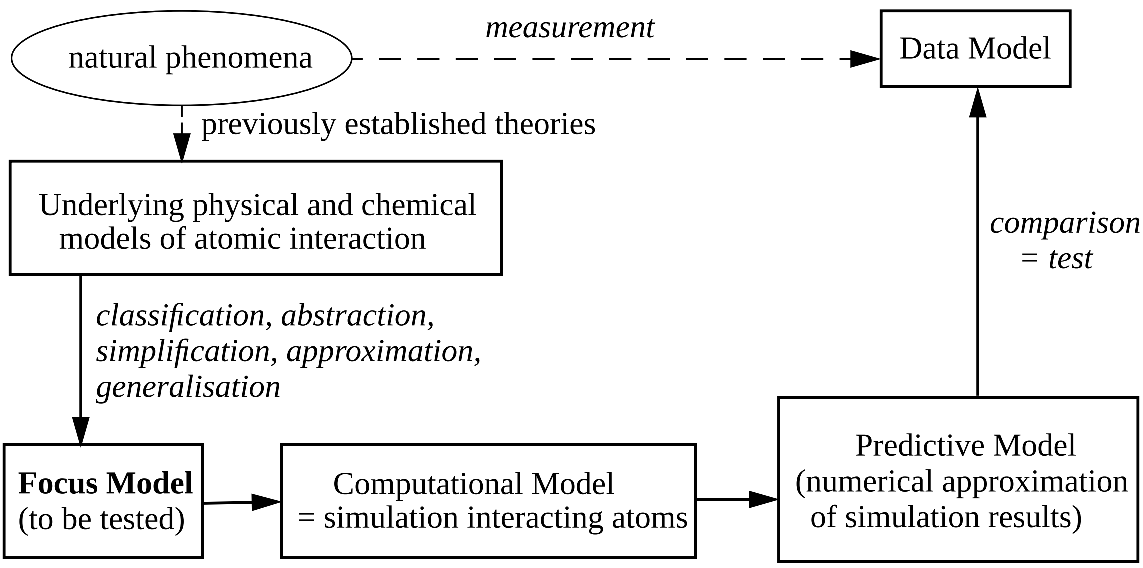

Figure 6. The models used in chemical simulation (adapted from [201])

Giere [176] goes further than a mere hierarchy of models, and also posits the existence of 'horizontal', 'vertical' and local 'radial' links between models. In this picture a theory consists of a whole structured cluster of objects: laws, models and analogies. This is surely the way forward: some of the confusion involving models in the past does seem to be due to the fact that many of the entities discussed as if they were a unitary structure are actually complicated composites with a series of different levels and processes.

This was foreshadowed by Suppes when he said in 1960:

> *"… a whole hierarchy of models stands between the model of the basic theory and the complete experimental evidence. Moreover for each level of the hierarchy there is a theory in its own right. Theories at one level is given empirical meaning by making formal connections with theory at a lower level." [431] p.260.*

2.3 Parts of the Modelling Apparatus

Given the above variety of definitions and deployment of models, I will attempt to ensure the clarity of terminology throughout this thesis by presenting a model of the modelling apparatus and then labelling the various parts. This will combine some of the best aspects of the philosophical analysis outlined above with a clear and unambiguous method of referring to the relevant parts and processes.

This picture of modelling broadly follows the model-theoretic approach, particularly as described by DaCosta and French in [129]. At its core is a state-space (following the approach of Beth and Weyl), which can be thought of as a systematic enumeration of all the possible data combinations that could result from the phenomena under study. So for the gas laws one might have a three dimensional space defined by coordinates for pressure, volume and temperature.

Different aspects of the modelling situation are all mapped into this one space. Thus the model description defines a subspace of these possible states, and it is this subspace which represents the model semantics. The data will be a series of points in the space and (in control situations there may also be a subspace representing the goals). The whole space can be circumscribed by relevant prior knowledge of the relevant possibilities (in the gas law case described above we know that we do not have to include temperatures below absolute zero). I will call the combination of model specification, model content and the relation between the two, the 'model-apparatus'. This is illustrated in figure 7 (the situation is shown in a simple way for presentational reasons, in more realistic cases the model content might not be continuous, connected, have a ‘hard’ edge (or membership function) or even be dense.

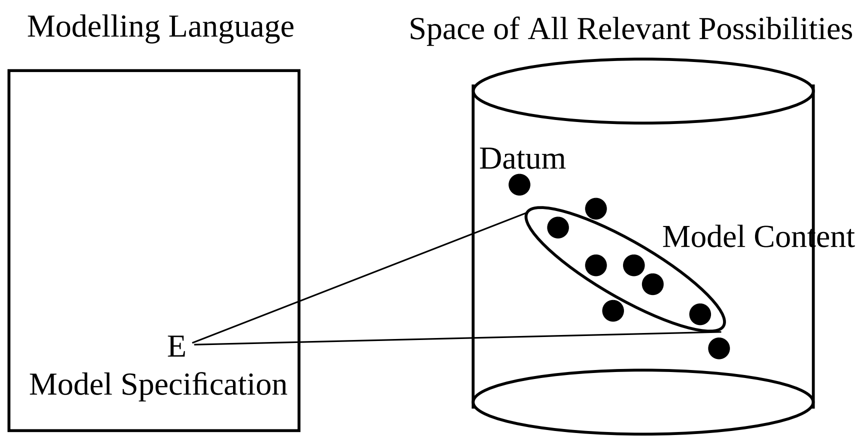

Figure 7. Model specification and semantics

The ‘modelling language’ is the set of possible specifications along with some relations between the specifications (equality, implication, etc.) and the way that these specifications are mapped onto the space of possible states. Following the model-theoretic approach the model content is the set of states which *satisfy* the specification. This modelling language could be a formal language, a subset of natural language, or even a real valued vector space. A good modelling language will not only be expressive enough to clearly specify the required possibility spaces, but also it will have explicitly defined relations that systematically reflect the corresponding relations between the possibility spaces. So, for example, the set-theoretic union of the possibility spaces might correspond to a disjunction in the modelling space. It is the fact that a good and expressive modelling language can accurately *reflect* relationships between the content of models that allows there to be syntactic as well as semantic accounts of modelling, but it is the non-syntactic approximation of models by their content (as described in [372]) which suggests that an account which includes the semantic side might be the more general.

The set of data (which may not be points but small volumes or even fuzzy areas) form what I will call the ‘data model’. This data model is itself a specification of the phenomena in ways directly analogous to the relation between model specification and the possibility space. In fact measure theory is the formalisation of such a relation, where the ‘modelling language’ is a language of numbers, and the content is the phenomena being measured, such that certain relations between the phenomena are homomorphicaly reflected in a relation on the set of numbers (e.g. concatenation and order).

The key idea (loosely following the approaches of DaCosta and French [129] and Redhead [372]) is that models are related and compared in the state space rather than via the model specification. A simple example of such a comparison is a measure of the error (for example the Root Mean Squared Error), which formalises the distance between the data model and the model content. In general a number of ways of judging the similarity and ‘distance’ between models is used. So, for example a phenomenological model content may *approximate* to a subspace of the theoretical laws’ content space, which may or may not be obvious from the specification of the phenomenological model and theoretical law. In other words it may not correspond to a consequence relation on the model specifications.

This picture suggests that there are at least three basic ways of relating model-structures: *firstly*, where one model content is contained in another (which I will call ‘subsumption’); *secondly*, where one model content is approximated by the model content of another (which I will call ‘approximation’); and *thirdly*, where the specification of one model is *described* by the specification of a second (which I will call ‘chaining’).

The subsumption of models, may be achieved by merely fixing some of the parameters of a more general model or law so that the restricted model only encapsulates some of the possibilities. For example the content of a general model of a damped and driven pendulum includes that of a free and frictionless pendulum. In general, the two model descriptions may be in different modelling languages, but mappable into the same state space. If the model specifications can be translated into the same modelling language and the inference relation on this language is complete (in the sense that every logical consequence can be proved using the inference relation), then this subsumption may be demonstrable as an implication in the modelling language. In general, however, this need not be the case. Relating models using subsumption is illustrated in figure 8.

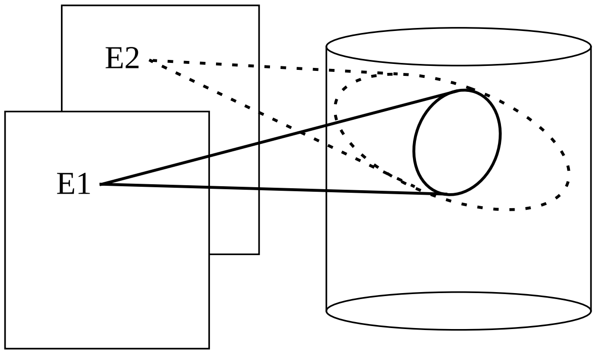

Figure 8. Relating model-structures via the 'subsumption' of model contents

The approximation of model contents can be achieved in a variety of ways. One may attempt it by careful manipulation of the model specification, although one has to be careful about this since a small change in the specification can cause a large change in the model content. Alternatively numerical simulation of the results of one model can be used as data to fit the second model. An example of this can be seen in the process of chemical simulation mentioned above, where the results of the simulation are fitted to a numeric function – in this case it is the role of the computer simulation to enable the approximate inference of a model where the analytics are too onerous. Relating models by the approximation of their content is illustrated in figure 9. This is comparable to Redhead's figure on page 151 of [372] in terms of relating model content. Such approximation of model content is implicit in the following quote from Cartwright:

> *"...we lay out a model and within the model we 'derive' various laws which match more or less well with bits of phenomenal behaviour." [84] p.161*

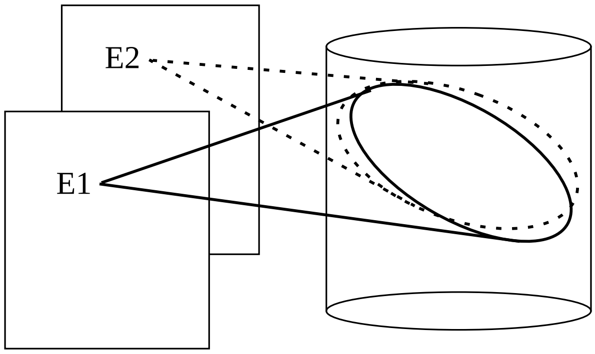

Figure 9. Relating model-structures via approximation of model contents

The relation of models by chaining is where the specification of one model-structure is used as the content for another. A example of this is how the data model specification (typically a list of N-tuples) is approximated by the content of a phenomenological model. This thesis is going to present a characterisation where model specifications are modelled by numbers, so that these act as a useful guide as to their complexity – which means that ‘complexity’ would be a (numeric) model of models of data models of phenomena! Relating models by chaining is illustrated in figure 10.

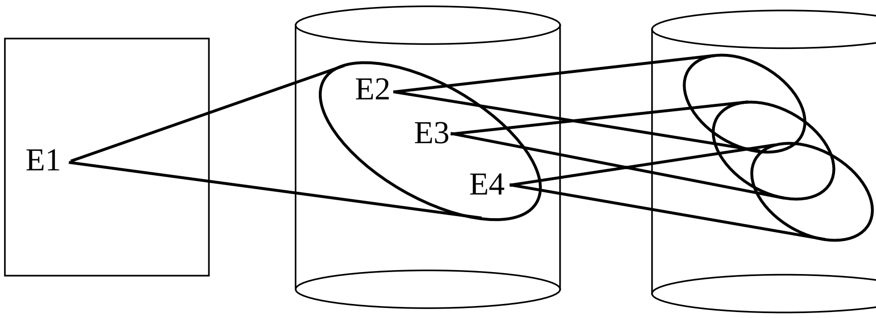

Figure 10. Relating model-structures via ‘chaining’

In the set-ups described above (section 2.2 on page 28), one arrow could actually represent a combination of several of these methods of model relation. The example from chemical simulation illustrates the use of all three of these types of model relation, as well as a combination of them: the data model and the predictive model are related by chaining; the predictive model is related to the computational model by approximation; the computational model is related to the focus model by subsumption (if it has been programmed accurately) and the focus model is related to the underlying theory via a combination of subsumption and approximation.

Note that I am not making a strong epistemic or realist claim here about what models and theories *are*, but merely using a simple model-theoretic structure as a building-block to model some of our constructs in an useful way. This is sufficient for my purpose here – I will leave the question of whether this reflects the actual use of models by scientists to others.

Different properties are usefully related to either the syntactic or semantic sides of the modelling apparatus. Thus the rate of error is a property of the model semantics whilst formal inference is more usually *attributed* to the model specification (even if, as a consequence relation, it can be *defined* using the content subspaces). The falsifiability of a model is more usefully associated with the model semantics: to take an extreme example, if we look at the null model (‘anything may happen’), this would correspond to the whole possibility space – in such a case there is no possible datum which would disconfirm the model. The falsifiability of a model (or *specificity*) may be associated with a ‘volume’ measure on the space *not* in the model content (for more on this and some of the consequences of this picture of modelling see [333]). These properties are summed up in table 1.

model syntax	model semantics
context-independent specification, formal inference, (complexity!)	error w.r.t. data, approximation, inclusion, specificity

Table 1: Some properties ascribed to the model syntax and semantics

My thesis will be that ‘complexity’ is a property that is more usefully attributed to the syntactic side of the modelling apparatus, because of the existence of simplification and the importance of the modelling language to practical judgements of complexity.

2.4 Models in Machine Learning

Another field where models have been actively considered in their own right is that of machine learning and artificial intelligence. Here a model can be any internal construct of an (artificial) agent that helps it decide what action to take, though usually the term also has predictive and representational connotations. These models will often be induced from the data, so there may not be an encompassing theory (other than what is implicit in its own structure or representational system). Thus the term has far more of an ad-hoc and pragmatic flavour here.

Thus in contrast to the top-down elaboration of the philosophical account of modelling described above, is the bottom-up process that frequently occurs in machine learning. Here the question is not ‘Why isn’t the general formal theory adequate on its own?’ but ‘Why isn’t the simple data model enough?”. A number of answers can be given:

- We know from experience that data sequences are rarely repeated exactly;

- We also know that, frequently, factors irrelevant to our purposes will have affected our data (e.g. ‘noise’);
- The data model does not allow us to predict in circumstances we have not come across before;
- We sometimes have prior knowledge that allows us to make a good guess at a more complete set of possibilities without trying them all (e.g. interpolation in continuous processes);
- The data model is not a compact form of representation;
- It is difficult to ascribe meaning to a set of raw data;
- Direct formal manipulation of the data is difficult[13].

All these lead in the same direction, namely towards some encapsulation as a specification in a more abstract (and maybe more general) model description via its (possibly approximate) encoding in a modelling language. The idea is to find a model description in a modelling language whose content closely approximates the data using what prior knowledge one has of the source of the data. I will claim that the *difficulty* in finding such a model specification is a useful characterisation of its complexity.

Given the nature of data gained by measurement from natural phenomena and the nature of descriptions in formal modelling languages, it would be very surprising if a compact or convenient description were found which exactly specified the data model – but, more fundamentally, it would not be very helpful in terms of generalising to areas of the possibility space were no data had been collected. Thus it is essential that the modelling language (and search process) is such that it does *not* result in a specification that directly describes the data. Almost always a model specification is settled on which either covers (i.e. includes) all the data points[14] or else approximates them to within an acceptable degree of error.

This is particularly sharply illustrated in the ‘over-fitting’ of neural networks, that is when the output of the network mimics the training data too closely and as a result (it is sometimes claimed) it does not perform well on new data from the same source. The

13.Though this particular difficulty has been significantly lessened with the advent of cheap computational power.

14.For example, in econometrics, a random distribution term is built into the model, so that the model will cover the data.

response is often to alter the structure of the network so that the space of functions that the network can produce is restricted, so the network is forced to abstract from the training data [401]. Of course, it is critical to the above strategy that the restriction utilises some prior and relevant domain knowledge concerning the nature of the data source, so that the network is restricted in an advantageous way, but this is often only specified implicitly (e.g. in continuity assumptions).

The last example highlights the importance of the modelling language to the modelling process. This is sometimes called the ‘bias’ in the machine language literature. Although the search for appropriate models is usually characterised as a process acting within a fixed modelling language, a more general picture is when both the modelling language and the choice of models are variable. Usually a change in the underlying modelling language is a more radical step, as such it has been associated with a change in paradigm as portrayed by Kuhn [274]. Toulmin associates it with the major advances in science:

> “*The heart of all major discoveries in the physical sciences is the discovery of novel methods of representation, and so of fresh techniques by which inferences can be drawn – and draw in ways which fit the phenomena under investigation. The models we use in physical theories, … are of value to physicists primarily as ways of interpreting these inferring techniques, and so of putting flesh on the mathematical skeleton.” [438] p.34*

Later I will argue that a change in the modelling language is one of the most powerful ways in which simplification is achieved (section 5.7.4 on page 123). In machine learning the importance of the choice of bias and the possibility of dynamically altering it to improve the efficiency of the search process has been explicitly discussed (e.g. Russel and Subramanian in [393]).

Frequently a model is used in the active sense[15] that is it is used to predict some consequences given some initial conditions. So the model specification is converted into a process (this is clearest in the case of a computational model, when the model is encoded as a program and then run). This is particularly obvious in model-based control mechanisms where the error in the prediction of the model is used to adapt the model.

15.But not always, for instance a fashion icon may be taken as a model of fashion in a whole subsection of society.

To reflect the fact that, in use, models are dynamic objects, some authors prefer to talk about the ‘modelling relation’ rather than models (notably Rosen [385]). Here a formal inferential process is related to a natural process by encoding and decoding relations. The action of a natural process is modelled by encoding the initial conditions into the formal model, from which an answer is inferred by some symbolic process, which is finally interpreted back as a prediction about the result of the process. This can be seen as a dynamic version of Hughes’s DDI (Denotation, Demonstration, Interpretation) account of a model [236]. The modelling relation is illustrated in figure 11. In terms of category theory the formal process is said to model the natural process if, given the encoding and decoding mappings, the diagram commutes.

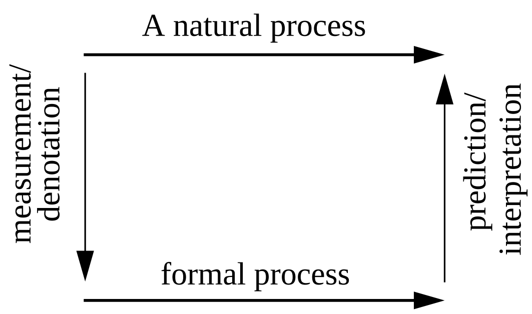

Figure 11. The ‘modelling relation’ from [385]

The connection between the former, static account of modelling and this dynamic account is the animation of declarative knowledge via the use of an inferential mechanism. Thus an equation is used to predict outcomes, by the substitution of variables and its solution, a program is animated by the program interpreter or processor. There are essentially process models without an explicit declarative counter-part (like a trained neural-network), but declarative model specifications tend to be more common. The prevalence of the declarative representation of knowledge is due to its flexibility: for example the same ideal gas equation can be used to either predict the pressure from the temperature and volume or the volume from the temperature and pressure.

2.5 The Philosophical Background to the Rest of this Thesis

My overall stance, from which I approach this thesis, is close to that of Giere[16] [176], except that I am not aiming to model science. My stance might be called a pragmatic modelling stance in that what I am aiming for is a *useful model* of complexity. So, for example, when I am arguing what the property of complexity can be usefully attributed to (in section 3.3 on page 47), what is most important to me is to identify the most useful object for that attribution. I will argue that it is most useful to attach complexity primarily to the model specification and then allow for its post-hoc projection onto the model content and the target phenomena when the mappings to these are sufficiently constrained. These constraints could come from the nature of phenomena but could equally come from elsewhere, for example in the established methodology of a field. As Cartwright puts it:

> *"It is precisely the existence of relatively few bridge principles that makes possible the construction, evaluation and elimination of models … it strongly increases the likelihood that there will be literally incompatible models that all fit the facts so far as the bridge principles can discriminate." [84] p.144*

My view on the constructivist-realist divide is that our models are at least somewhat constrained by the natural world, but not completely so. In some domains we may be more constrained by the world than in others: our model of a salt crystal as a cube may well reflect the world; our linguistic models of natural language may inevitably be socially constructed to some degree. As Giere puts it:

> *"The real issue, as I see it, is the extent to which, and by what means, nature constrains scientific theorizing."* [176] *p. 55-56*

Although the arguments herein follow more easily from a constructivist standpoint, they apply also to modelling as envisioned from a realist perspective.

To sum up: the aim of this thesis is to produce a general and useful *model* of the notion of complexity. The various existing formulations of complexity along with the other examples that I introduce play the part of the data. The main work is thus to find the model that matches these as far as possible but is also as 'useful' as possible. This will involve the formulation of the best 'modelling language' for its expression (i.e. one that matches prior reasoning about the domain's nature), so that when we have finished our search process we end up with a specification of the concept that is not too complex but

---

16.Griere calls his position "naturalistic realism".

adequately approximates our ‘data’. In this regard, merely saying that every situation has its own complexity would be like sticking to the data model in the above, and this would be an inadequate account if we are ever to apply our account to new situations. As Einstein is supposed to have said when asked how complex a theory should be:

> *“Everything should be made as simple as possible, but not simpler.”*