This is the final submission of

Ruttkamp, E.B. 2006. "Philosophy of Science: Interfaces between Logic and Knowledge Representation", *South African Journal of Philosophy*, 25(4), 275-289.

PHILOSOPHY OF SCIENCE: INTERFACES BETWEEN LOGIC AND KNOWLEDGE REPRESENTATION

Emma Ruttkamp

Department of Philosophy

University of South Africa

PO Box 524

Auckland Park

2006

E-mail: emmar@uj.ac.za

PHILOSOPHY OF SCIENCE: INTERFACES BETWEEN LOGIC AND KNOWLEDGE REPRESENTATION[1]

ABSTRACT

Against the background of a discussion of knowledge representation and its tools, I offer here an overview of my research in philosophy of science. I defend a relational model-theoretic *realism* as being the appropriate meta-stance most congruent with the model-theoretic view of science as a form of human engagement with the world. Making use of logics with preferential semantics within a model-theoretic paradigm I give an account of science as process and product. I demonstrate the power of the full–blown employment of this paradigm in the philosophy of science by discussing the main applications of model-theoretic realism to traditional problems in philosophy of science.

I discuss my views of the nature of logic, and of the role it has in philosophy of science today. I also specifically offer a brief discussion on the future of cognitive philosophy in South Africa. My conclusion is a general look at the nature of philosophical inquiry and its significance for philosophers today.

INTRODUCTION[2]

Balancing on a branch in the tree at the end of our garden, clutching Aristotle's *Posterior Analysis* in one hand, I was, at the age of about 15, contemplating the seductive Greek images of truth, precision, equilibrium, knowledge, and of course, wisdom. On my – almost daily – excursions to my beloved tree I used for years to intrigue myself with thoughts of finding harmony in the universe, seeing symmetrical patterns of human thought impressed in the world around me, and always finding everything – the universe, my world, the world – resting safely and reassuringly on steady strong pillars of wisdom and truth.

It is quite likely that I was to most people an insufferable little brat who thought she knew everything, but that never bothered me much, I had this idea – due to the Greeks I was reading - that by doing gymnastics or ballet, studying mathematics, and listening to classical music, I would be able always to see the very structure of the world and know exactly what was happening – I would

---

[1] Inaugural lecture read at the University of Johannesburg, 24 May 2006.

[2] I am deeply grateful to colleagues in my field who have played, one way or the other, a significant role in my research over the past 10 years. I especially mention here Johannes Heidema at Mathematical Sciences, Unisa; Willem Labuschagne at the University of Otago in Dunedin, Gabor Kutrovatz and Laszlo Ropolyi at Eötvos University in Budapest, Gabor Zemplen at the Max Planck Institut für Wissenschaftsgeschichte in Berlin, Jan Wolenski at the University of Krakow, David Spurrett at the University of Kwa-Zulu Natal, Don Ross at the University of Alabama at Birmingham.

be pursuing the beauty of pure forms, discoursing in the language of the universe, and discovering harmony everywhere after all, how else could it turn out?

Needless to say, obviously totally chaotic, and mostly without even a smudge of pure precision, elegant truth, or stable wisdom in sight! How could this be? The stage was set, the tools were at hand … what went wrong? Perhaps the short answer lies in the mundane fact that I, after all, turned out not to be a Greek goddess, but rather just Emma. The long answer lies in the intricacies of knowledge representation. I shall, rather than bore you with the short answer, for the rest of my time with you tonight, try to tease out some of the complexities contained in the long answer. And I shall do this by picking out from the myriad relations between logic and scientific knowledge the second-order interface between logic and knowledge representation.

LOGIC[3]

Logic can be helpful in two obvious ways: It guides intuition about at least the respectability of arguments via informal logic, and it offers certainty regarding validity of arguments courtesy of formal logic.

Informal logic can save us from believing arguments like: 'Everyone has a right to education. Therefore the government should offer excellent tuition to everyone with no obligation (in terms of school fees, or anything else) on the part of learners'. Well, this sounds logical does it not? Why would one want to be 'rescued' from an argument like this? Because it is *not* logically valid. Unfortunately there are many situations in which the premise 'Everyone has a right to education' is true, while, at the same time, the conclusion is false, and thus, there are many cases in which this argument is invalid. Certainly my everyday world is an interpretation (or valuation, or possible world) in which this argument is invalid. So, informal logic may teach us to never rush into interpretations, but rather to tread carefully and lightly with the angels.

Formal logic falls into studies of formalisation of natural language (such as Zulu, English, Greek, etc.) and human (or automated) theorem proving. The former uses mathematical symbols to express knowledge claims in terms of well-formed formulas, the latter uses rules of inference to make formal deductions from sets of axioms. Here I concentrate on the formalisation of natural language, and thus on knowledge representation.

[3] I am grateful to my friend and colleague, Willem Labuschagne, currently at the Department of Computer Science, Otago University in Dunedin, New Zealand, for the direction some of his earlier work at Unisa, and his own inaugural lecture (Labuschagne 1994) gave the final version of this lecture. Much of the content of the first sections of this lecture is owed to his writing, combined with the influence of Johannes Heidema from Mathematics at Unisa.

You could very well ask why – apart from adolescent dreams of pure truth – would anyone want to use mathematical symbols to express their thoughts? Weirdly enough, the answer actually is not too far removed from these dreams after all: Mathematical symbols offer a mechanism for

- Avoiding fallacies, ambiguity, and the overall 'meaning-loadedness' of natural language
- Establishing clear similarities between two different situations (or systems) by describing both in a standard form of notation
- Keeping the distinction between meta- and object levels firmly in place (this is essential for any talks about truth to be meaningful, as we shall see later)
- Applying a theorem-proving algorithm to explore the consequences of knowledge claims.

KNOWLEDGE REPRESENTATION

1. Background

Talk of knowledge representation refers to the act of describing some system of interest with a formal (or 'artificial' as opposed to 'natural') language.[4] As part of my research for my DLitt et Phil degree, I investigated the application of basic principles of knowledge representation to traditional philosophical descriptions and explanations of the processes of science, because I suspected these principles of offering an interesting way to rescue science both from barren instrumentalists and deranged post-modernists.

From a logician's point of view, knowledge representation includes the following: A specific system in reality that is focused on, a conceptual model of this system, and an application of axiomatic method. I have claimed that the same three elements can also be identified in second-order analyses of the processes of science (see Ruttkamp 2002). Let us now briefly consider each of these elements of knowledge representation in turn.

1.1 The 'intended' system

In terms of knowledge representation, focusing on a particular real system, referred to as the 'intended' system, serves mainly a descriptive function. Here the choice of system is determined by the function and objectives of the program being developed. Examples of such systems could be a room to clean, a motorboat, the set of natural numbers, the gas molecules in a flask, a light-fan system, and many others. In science the choice of system is determined by the disciplinary matrix of the scientists in question, and other context specific factors, such as research intentions and goal,

[4] One field in which knowledge representation obviously is extremely important is that of Artificial Intelligence. This is so, because artificial agents, or robots (whatever form they may take), are given information in terms of programs (which are strings from some formal language).

field of application, and others.[5] It is important to note that it is always a *slice* of reality that is investigated, described, or examined and never somehow reality in all its multiplicity.

1.2 The conceptual model

As part of the mechanism of knowledge representation a conceptual model is always formed of the intended system in question. This 'model' can be a representation in one's mind (or brain – suffice it to say it can be 'mental') or it can be expressed in some linguistic form. It is 'suitably' idealized, generalized, or simply skewed in terms of being an abstraction of the intended model, a representation representing only features of the intended system relevant for intended research, and then mostly in general terms. A conceptual model never represents the intended system in its full glorious actual presentation.

It includes a certain set of objects, and certain properties of these objects and relations between them, representing the way things are, or work, in the intended system. Say our intended system is a simple light-fan system. Let us agree to formulate the set of objects in the conceptual model of this system as consisting of 2 elements: the light, and the fan. Let us note properties such as the fact that either the fan or the light can be switched on or off at particular intervals.

1.3 The axiomatic method

The assumptions made about the intended system must now be made explicit. This is an idea that originates with Euclid's writing down of the first comprehensive exposition of geometry. At this stage of knowledge representation, a language is formulated or 'designed' strictly guided by the relevant conceptual model.

Say we design a language L containing a set of individual constants, called Ind. with its elements being a name for each object of relevance in the intended system, say in our case, Ind. = {l, a}, where 'l' stands for 'light', and 'a' denotes the 'fan. L also includes a set of predicate symbols, describing properties of and relations between the objects described in Ind. Let Pred = {(IsOn, 1)}. (We could include predicate symbols with an arity of more than one, but I am keeping things simple. L could also include a set Fun., of functions symbols, and a set Var., of variables.)

Part of specifying our language is to specify the logical connectives that will be used and defined in the syntax of the language, say at least **5**, **ω**, **ϖ**, and **6**. (The universal and existential quantifiers will also be defined in the case of a quantified language, but in the current case this is not applicable.)

---

[5] See Holston's definition of 'themata' in Holston (1995).

The semantics of the connectives may be specified by truth tables or by logical axioms. Then other extra-logical axioms regarding the intended system may be formulated, specifying how the system works, for instance, IsOn(l) 6 IsOn(a).

It is important to realize throughout this discussion that actions here are guided by the conceptual model of the intended system, and *not* by the actual real system. It is impossible to include anything into the axiomatic descriptions which was not built into the conceptual model. Notice the use of the term ‘build’ here, the conceptual model is something that is *created* in some mental space by the observer-programmer-scientist interested in the relevant real system at the time.

There are also other factors we have to take account of that affect the accuracy of the axiomatic description:

- The ‘grain size’[6], i.e. the decision of what to take as the ‘building blocks’ of the system
- Preference (and intention)
- Ignorance

These points emphasize the fact that the conceptual model of the intended system is never ‘perfectly accurate’.

2. Meaning and reference

Now, after the axiomatic method has been applied, questions can be asked about the intended system, and it can be checked if relevant propositions follow logically from the axioms.[7]

In terms of scientific method, all of this is rather Aristotelian of course. Aristotle’s method of induction-deduction implied that science is a progression from observations to general principles (inductive stage of scientific inquiry), and from general principles back to observations (deductive stage of scientific inquiry). In other words, scientists induce explanatory principles from the phenomena to be explained, and then deduce statements about the phenomena from premises which include these principles. Notice also the interesting fact that it was Aristotle who shifted the emphasis in philosophy from the Socratic examination of the nature of knowledge to the problem of representing knowledge (see Sowa 2000).

[6] Labuschagne’s term.
[7] This is how ‘extensions’ of sets of well-formed formula are formulated, given that extensions are the relevant axioms together with the logical consequences that may be deduced from them.

Now, during this (deductive) stage of events, the question thus becomes how to link the conceptual model 'back' to the real system it originated from, i.e. how to give 'meaning' in terms of some real system, to the content of the conceptual model. In terms of knowledge representation this is done by formulating an interpretation for the language in question. This has proven to be a task whose success (in terms of 'fit' between real system and conceptual model) cannot always be guaranteed. This is the case for two reasons: First, it is, after all, an idealization being linked to a particular, detailed system, and that implies at least some degrees of skewed mutual fitting (except if one is Aristotelian enough in one's beliefs that the general exists only in the particular not to be bothered by this); and second, axioms are wffs in a formal language, and this language can be interpreted in many different ways (Tarski 1935, 1982).[8]

What is an interpretation? An interpretation is a mathematical structure. More specifically, it is an ordered pair (D,f) where D is a domain of objects, (a domain of discourse, or universe of discourse) and f is a so-called denotation function telling us three things:

- ❖ Which object in domain D is denoted by each individual constant
- ❖ Which function in D is denoted by each function letter in Fun (not applicable in our current example)
- ❖ Which property or relation is denoted by each predicate symbol.

Typically, there are as many interpretations as there are pairs (*D*,*f*). Some of these interpretations will be models of the relevant set of axioms, and among these may be the conceptual model of the intended system. Notice here that a model in the Tarskian sense is also a mathematical structure, and specifically it is an interpretation of the relevant language that makes the chosen axioms of the language true. Note also that 'true' here means that denotations of individual constants and function symbols 'fall' within the denotation of the predicate symbols, which are always subsets of the domain *D*. Truth is thus a relation between sentences and interpretations.

In our example above, an interpretation for L could be I = (D, f) where D is the light-fan system depicted by our language, and is written D = {l, a}. Then we can choose the denotation function f as mapping each element of Ind onto itself in D[9]; i.e. f(l) = l, f(a) = a. Furthermore we choose the

[8] By the nineteen fifties Alfred Tarski's model theory, and his views in particular on truth and logical consequence (see Tarski 1935, 1982), had matured into a definition of truth in terms of relations between sentences and interpretations. His studies of the properties between sets of sentences and classes of interpretations opened up new horizons for studies in formal logic in general, and knowledge representation in particular. This is important because in this context truth is presented in a relational way, rather than as an absolute property of sentences.

[9] Such an interpretation is called a 'term-interpretation'.

denotation of the predicate symbol in Pred as follows: f(IsOn,1) = {a}. This interpretation is not a model of the sentence IsOn(l), since this sentence is false under this interpretation.

There are some measures that can be taken (such as ensuring the completeness of the set of axioms)[10] to try to limit the number of models of a set of axioms, but they are not always indicated.[11] Keep in mind also though, that it might not always be desirable to have only one model of the set of axioms, but that rather a whole class of models might be useful. In such cases each model represents a state of the system (at a given time).

3. Truth

The most familiar, common-sense definition of truth is that it is a property of sentences that correspond with reality. In other words, to determine truth, one looks at some *unique* interpretation, namely the real world, and checks whether the sentence adequately describes the state of affairs.

The problem is that humans interpret and process 'data' from reality by using each their own internal mental models of the world. Thus the same object or event in reality may have more than one 'meaning' or interpretation if interpreted by different people or at different times. Perhaps, then, you might think, a designated refined version of reality should be used as the touchstone of truth; let's say the scientific model of the world? The point is that this model is still being built, and moreover, it will always be subject to revision at various times. Worse, it might be that different possible worlds might be equally well supported by science at a given time.[12]

Actually, science mostly attempts to build a class of models rather than a single one, because it has to allow for things that have not yet happened, or might never happen. Hypothetical states of the system must be considered for obvious reasons, often obviously to do with safety, but also with explanation, and predictability. Specifically, counterfactual statements play an important part in science. For example, consider the following statement: 'If you did not eat the chicken, you would

---

[10] For the axioms to exclude all but the conceptual model, the set of axioms must be at least complete – i.e. their extension consisting of the axioms and their consequences must be such that if any sentence (closed wf) is picked from the language in which the axioms are formulated, then either that sentence belongs to the extension, or its negation does. If the axioms are not complete, i.e. if there is a closed wf such that neither the wf, nor its negation is deducible from the set of axioms, then the set of axioms will have at least two models, one in which the wf is true, and one in which its negation is true. (This follows from Gödel's theorem that every consistent extension has a model.)

[11] Moreover, as far as more complex systems go, Gödel's Incompleteness theorem actually confirms that it is impossible to arrive at a complete axiomatisation of some of these systems by simply writing down a list of axioms representing what we know.

[12] Typically (traditionally) the idea was that scientific statements that 'fit' *all* possible worlds are the worthwhile ones, while statements false in one or more possible worlds were regarded with mistrust. On the other hand, a more human and more accurate view of science as a body of revisable knowledge claims is compatible with the model-theoretic fact that the more knowledge an agent has, the fewer models she has, and that a claim with only *one* model carries much more information than a barren tautology.

have been well'. You did eat the chicken, so the antecedent of this sentence is false in reality, but it manages to place one in a different possible world (other than the 'real' one) in which you are well and not slithering around in pain. Counterfactual reasoning, as opposed to reasoning directly about the 'real' state of a system, helps increase understanding of some important aspects of the current state of affairs.

This is why it is so helpful to have a mechanism by which it is possible to work with many possible worlds, which, in its turn, is why knowledge representation becomes interesting in second-order discussions of science: a typical formal language presents us with infinitely many interpretations, but formal logic presents us with many different tools to trace relations between languages and their interpretations. This is also why it is nonsensical to refer to a sentence as 'true' when it describes the state of affairs in 'the real world'. Sentences are true *relative* to some interpretation (or mental model or possible world).

IMPLICATIONS FOR THE UNDERSTANDING OF SCIENCE AS HIGHLIGHTED BY A KNOWLEDGE REPRESENTATION PARADIGM

1. The tentativity of scientific knowledge

I view science as an historical process — or even a set of historical processes — which continuously, in many different complex contextual ways, revises itself. I believe that to stay in the game called philosophy of science we need to find ways in which we can magnify the infinitely many different pathways existing among theories, their models, and real systems during such processes. My view regarding the nature of scientific knowledge fits well with the characterisation of science and its method(s) offered by the Finnish logician and philosopher Illka Niiniluoto (1999: 5). He (ibid.) concludes that science, if successful, 'will ... have tentative results, in principle always open to further challenge by later investigations, which constitute what is usually referred to as the 'scientific knowledge' of the day'. Scientific knowledge, in these terms, is revisable and may be represented in many different ways, but it is, I believe, also continuous, sometimes cumulative, and always rational.

After positivism, philosophy of science was in crisis. The way out seemed to consist of two available options (Niiniluoto 1999: 14): A choice either for following a descriptive historical method, building on Kuhn and Lakatos's work (and perhaps ending up in empty relativism), or a choice for continuing with the positivist quest for logical and quantitative perfection in philosophy of science (and ending up on equally infertile ground). I believe there is a middle way, a *relational* way, rather than a relativist way, that does not choose any of the available sides. As Niiniluoto

(1999) also stresses, a strict distinction between historical and formal methods in philosophy of science is unwise. On the one hand, we need logic (as I hope I have given you a glimpse of), and there are many other applications, for instance in verisimilitude studies, of non-classical logics that strengthen the view that logic today still has much to offer philosophy of science. We also, on the other hand, need philosophy to formulate definitions for terms such as truth and justification, progress and realism. Science is a rational cognitive but incomplete enterprise (ibid.). We 'live in ongoing stories', even in science (Lötter 1994: 157), but, at the same time, philosophy of science 'cannot leave science as it is' (Niiniluoto 1999: 14). Normative questions about scientific enquiry and knowledge have to be asked and answered, the social role of the processes of science is in need of philosophical reflection, and even the ethics of science is in constant need of philosophical attention.

2. The structure of scientific theories

I have formulated a model-theoretic depiction of science (Ruttkamp 2002) as a midway stance between the so-called 'statement' and 'semantic' views of theories, which are the two traditional schools of thought in studies of the structure of theories. Given its mathematical logic basis, my model-theoretic realism retains the positivist (statement) notion of a scientific theory as a (deductively closed) set of sentences expressed in some appropriate classical first-order language, but also acknowledges the post-positivist (non-statement, or 'semantic') emphasis on the interpretative and referential role of the conceptual (i.a. mathematical) models of these theories. There are many reasons for such a middle-of-the-way or mixed approach: (These reasons are explored in depth in Ruttkamp (2002)).

Briefly, the 'theory' is needed to avoid empty relativism. The theory is the one specification which all its models and their substructures have in common, and I need to retain this relationship between theories and models, given that my account is *relational* rather than *relativist*. At the same time though, models or mathematical structures are needed, because the process of science manifests itself in a differentiated way. This reflects the fact that the actual growth and internal turmoil theories undergo before they may be applied, revised or, ultimately, succeeded, take place at the level of mathematical structures interpreting the theory.

The above points are important and have indeed been acknowledged by some philosophers working in this field (see Ruttkamp (2002) for detailed discussions of some of the most important of these). I have though come to realize only recently that the most important need for models and more specifically for a truly model-and-theory analysis of science is presented by the fact that contextual factors — so-called 'themata' (Holton 1995) — and the factors causing 'preference' of one application or model of a theory above another are also changeable and context-specific, and above all, not part of the syntax of the theory. These factors are better captured in a preferential framework, where models can be ranked according to preference determined by heuristic contextual factors. I thus offer a model-theoretic account of science, based not on syntactic inference rules, but on a preferential semantics. More about this later.

3. Abstraction and idealisation vs 'concretisation'[13]

The 'idealised' nature of science, the fact that the laws of science 'lie' as Nancy Cartwright (1983) famously pointed out, and the need for boundary conditions and *ceteris paribus* clauses, have been part of doing science and philosophy from ancient times, starting at least with the struggle between the Pythagoreans and supporters of the method of saving the appearances. The need to abstract relevant and general information from a particular real system at a certain time has already been discussed in terms of knowledge representation above. The complexity of the process of 'linking' linguistic expressions of knowledge 'back' to real systems was also touched on, and will be discussed in more detail below under 'Realism'. The point is, in other words, that a model-theoretic depiction of science can meaningfully engage, even if only in a clarifying sense, with these difficult and many faceted features of science.

4. Realism

Linking 'back' to real systems from interpretations (conceptual models), is in model-theoretic terms a process that includes also two other kinds of mathematical structure – namely empirical reducts and empirical models[14] – as well as myriad links among these different kinds of structure. Recall that mathematical model theory and its definitions of interpretations of (sentences in) formal languages present the possibility of *many different* models of a given theory T (in language L). These models are interpretations of T's language such that a model of a theory sees to it that every predicate of the language of the theory has a definitive extension in the underlying domain of the model which satisfies T in the Tarskian sense.

Please see *Figure 1* for what follows.

[13] Cartwright (1989)'s term.
[14] Although adapted to fit my objectives here, the notion of an empirical substructure originates with van Fraassen (1980) and that of an empirical reduct with Kuipers (2000). See Ruttkamp (2002) for detail.

Figure 1: A model-theoretic account of science I

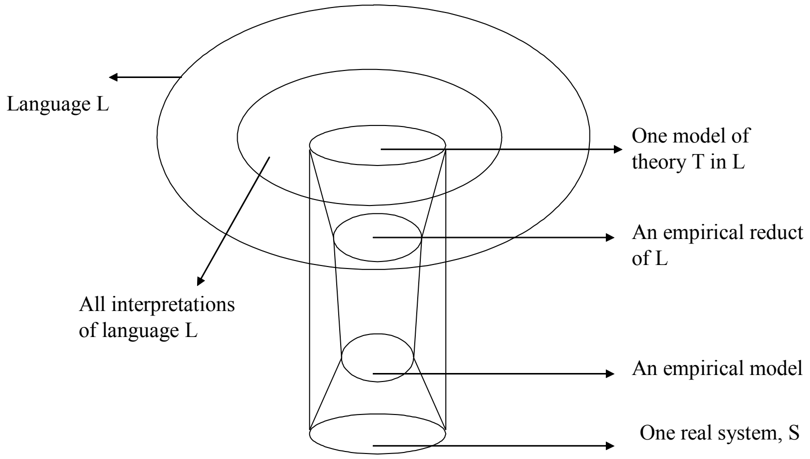

Focusing on a particular real system at issue in the context of applying a theory implies a specific empirical set-up in terms of the situationally observable properties and measurable quantities of that particular real system. To engage with this set-up, a relevant conceptual model of the theory, chosen for particular reasons at a given time, is trimmed ('reduced') in a specific way. In the context of science and theory application, and considering relations of reference, it makes sense to concentrate only on the predicates in the mathematical model of the theory under consideration that may be termed 'empirical' predicates (in the particular context of application). This is how an empirical substructure, or 'reduct' is formed. A 'reduct' in model-theoretic terms is created by leaving out of the language and its interpretations some of the relations and functions originally contained in these entities. This kind of structure thus has the same domain as the model in question, but contains only the extensions of the empirical predicates of the model. Notice that these extensions may be infinite since they still are the full extensions of the predicates in question.

Now, looking at theory application from the 'other' side, i.e. the side of real systems, from the experimental activities carried out in relation to the relevant real system, a conceptualisation of the results of these activities (such as performing certain observations and experimental tests), may be formulated. This (mathematical) conceptualisation of data may be represented as an 'empirical model'. This is a model then which has been determined by data gained as a result of interaction with some aspects of the real system in question (and not by the models or empirical reducts already constructed). The empirical model contains finite extensions of the empirical predicates at issue in the empirical reduct, since only a finite number of observations can be made at a certain time. In other words, an empirical model usually involves only finitely many elements from the domain of the model (or reduct) and from the extensions of the predicates.

In terms of 'reference' then, should we find that there is a one-to-one isomorphic embedding function (a relation of structural identity) from the empirical model into a certain empirical reduct, this would imply that there exists some relation of reference from our original theory via at least one model and at least one of its reducts, to the relevant real system (represented as at least one empirical model). This is close to what Bas van Fraassen (1980) termed 'empirical adequacy'.

But does this imply reference in a realist sense? Yes, it does. Under the interpretation chosen, the terms in the language refer to the interpretation's domain of discourse (representing an aspect of some real system in reality), and the interpretation (model) makes the sentences containing these terms true (i.e. the denotation function of the interpretation satisfies the language). Some of these terms are empirical in the current situation, and if completely unrelated experiment(s) is (are) done on the aspect of reality at issue and it is found that the results can be translated into the language, L, of the theory under consideration, and also moreover, an isomorphic relation existing between this empirical model and some empirical reduct of a model of the theory expressed in terms of L is found, a link – however 'weak' – between the terms in the language, and objects and relations in some real system, does exist.

To summarise: A model interprets all terms in the appropriate relevant language and satisfies the theory at issue. In the empirical reduct are interpreted only the terms called 'empirical' in the particular relevant context of application or empirical situation. Think of this reduct (substructure) of the interpretation (model) as representing the set of all atomic sentences expressible in the particular empirical terminology true in the model. An empirical model — still a mathematical structure — can be represented as a finite subset of these sentences, and contains empirical data formulated in the relevant language of the theory.

In my context it is simply incorrect to conclude from the 'no miracles' argument for realism that if we are realists, we should be realists about a world exactly like what science describes and be done with it. Rather, I agree that we should believe in a world rather like what science describes, but stress that it is no simple matter to say what aspects of science refer to the world, what refer to bits of our models, what are and are not equivalent, and so on. Also, most importantly, these 'bits' of science do not fit together without serious work being done. (See Ruttkamp 2002.)

5. Truth

There are many different models possible of one language – infinitely many to be exact. The main implications of this are in terms of the notions of truth and reference. Model-theoretically truth is a relation between sentences in formal languages and their interpretations. In this sense, reflections on whether we are examining a 'correct' or 'true' representation of reality remain, at the least, naive. The slogan of a model-theoretic realism is 'truth without universality'. 'Truth' is relative to specific models, and so questions of truth are settled by focussing on conditions of verification in the semantic sense of defining interpretations of the scientific language on specified domains of discourse. So, there are elements of conventionalism in a model-theoretic approach to theories, in the sense that 'truth' is something that we 'create' by our (heuristic or pragmatic) choices (of interpretation), and not something dictated to us by nature (or philosophy).

At the same time though, the definition of 'empirical truth' in terms of empirical adequacy and an isomorphic relation between empirical models and empirical reducts cannot be 'forced' or manipulated.[15] For one thing, the nature of the empirical model is fully 'determined' by the features of the relevant real system and experimental activities carried out in relation to it. Moreover, keep in mind that empirical reducts finally link empirical models to the theory, given that reducts are substructures of the models in which the sentences of the language of the theory are true. Here, keep in mind that the choice of conceptual model(s) is determined by research intentions and thematic preferences of the scientists applying the theory within some accepted Kuhnian disciplinary matrix, or 'against' some background meta-theory. Moreover, the choice of model(s) is also determined by the syntax of the theory — i.e. the models are interpretations of the *language* of the theory such that the theory is 'true' in, or 'satisfied' by, them.

6. Under-determination and over-determination

There may be situations where there is a relation of reference between one empirical model and many theories – so-called under-determination of theories by data – or, vice versa, where there is a relation of reference between many empirical models and one theory – so-called over-determination of theories by data (Ruttkamp 2005). Both forms of determination are meaningful in a model-theoretic context. In terms of under-determination one empirical model is 'linked' to more than one theory. When Duhem and Quine first discussed this feature of science, it was supposed to introduce at least serious problems for, if not the end, of realism. But model-theoretic realists are sophisticated realists who don't care about one-to-one correspondence, only about semantic

[15] This relates to discussions (Ruttkamp 2002) around Putnam's so-called paradox (Putnam 1983).

reference – model-theoretic satisfaction and empirical adequacy are enough for us because we can trace many-to-many links.

The case of over-determination is also not problematic for the above reasons. We do sometimes though, want to know how and why we choose applications of theories at specific times, and this is where preferential semantics comes in (see Ruttkamp 2002, 2005). In my research I have mainly used Yoav Shoham's (1988) model-theoretic non-monotonic logic, which offers a fairly simple way of ranking models, to analyse and trace motivations behind preferences for certain empirical models at certain times (Ruttkamp 2002).[16]

By making use of non-monotonic 'ordering' I have shown (Ruttkamp 2005) that ranking of empirical models, induces ranking of conceptual models, and also, ultimately, of theories. This enabled me to draw a picture of scientific progress that is continuous, albeit highly complex and with many different pathways among models, theories, and real systems. These pathways are like spider webs, they have an overall recognisable pattern, but the number of links, and the fineness of the thread differ from spider to spider.

7. The context of discovery

Karl Popper famously, by rejecting the inductive method of science, shoved the context of discovery out of reach of philosophy of science, and into the outstretched hands of psychology. A model-theoretic representational preferential analysis of science's processes snatches the whole intricate story of contextual motivation 'behind' the 'birth' of theories right back and re-establishes it firmly in the realm of philosophy where, as it happens, I believe, it belongs.

8. Empirical-theoretical distinctions

The fact that in my depiction of science the empirical-theoretical distinction is placed at the conceptual level, means that it is not a unique rigid distinction. The positivists fruitlessly searched high and low for a single definition of correspondence rules or bridge principles to represent a (unique) method of translating theoretical terms into empirical (observational) ones. In my view, Rudolf Carnap's (1966, 1975) notion of 'reduction sentences' comes closest to the supple contextual kind of translation offered in my scheme of things by the model-theoretic formulation of 'empirical reducts' of conceptual models.

[16] Note that this kind of non-monotonic logic is a very small part of a very complex set of non-monotonic logics (see Delgrande, Schaub, Tompits, & Wang (2004) for a detailed discussion of the field).

9. Reduction

I recently with Johannes Heidema (Ruttkamp & Heidema 2005) proposed a ‘non-traditional’, or non-classical generalization of Ernst Nagel’s (1949, 1961) positivist definition of theory reduction in philosophy of science. I have since (Ruttkamp 2006) also investigated the possibility of adapting this definition of so-called ‘defeasible’ reduction to the mind-body problem.

Our definition is formulated with the aid of non-monotonic logic, for the good reason that non-monotonic reasoning by definition is about conclusions being drawn from sets of premises or ‘evidence’ by methods that are not strictly speaking deductive – think of Sherlock Holmes’s methods of ‘deduction’ (see Makinson 2005: 1) – in the sense that they involve ‘presumption and conjecture and the ever-present possibility of going wrong. … conclusions may be withdrawn as more information comes to hand, and new ones may be advanced in their place. This does not necessarily mean that there was an error in the earlier reasoning. The previous inferences may still be recognized as the most reasonable ones to have been made with the information available’. Translated into mind-body language, this seems compatible with multiple realisability, and also perhaps with some forms of eliminative materialism.

In a philosophy of science context, having an alternative definition for reduction is necessary, because theories sometimes are reduced in various ways to various degrees as part of normal scientific processes. Here reduction in the defeasible sense is a practical device for studying the processes of science, since it is about highlighting different aspects of the same theory at different times of application, rather than about naive dreams concerning a metaphysical unity of science.

My claim in terms of philosophy of mind is much more tentative. It is not about establishing a metaphysics of mind, and not at all about determining one to one relations between terms in psychological theories and terms in neurological theories. Mostly what is important is that it seems as if relations between the mental and the physical – if any – are always naturally tentative, complex, and are at least one-to-many, if not perhaps many-to-many. In this context, it is interesting that defeasible reduction allows for different relations of reduction between the same two theories at different times, according to different rankings (in terms of some preference) of the models of the reducing theory. Also, the rules according to which rankings are done allow for tracing of motivations behind certain reductions at certain times.

I have recently started to wonder whether relations of reference may not contain more elements of reduction than realized in traditional accounts of realism. I am currently working on a paper considering this against the background of Patrick Suppes's idea of a hierarchy of models and theories existing between the experimental and the final theoretical levels (Suppes 2002), while I am also in this context considering some of Carnap's work on reduction sentences (Carnap 1966, 1975).

REFLECTION

At universities, courses in logic are today mostly found either in mathematics, philosophy, or computer science departments. Formal languages, deductive systems, and semantics are mathematical objects, and as such, a logician is interested in their mathematical powers and properties (Shapiro 2000). Philosophically though, logic is mainly the study of correct reasoning as an epistemic activity, and so questions on the philosophical relevance of the mathematical aspects of logic arise (ibid.). Many different opinions have been offered on this question over the years.

My view is that, in an Aristotelian sense, the importance of logic as a field of study in philosophy should not be underestimated. At the same time though, I do not want to exclude philosophy from benefiting from the mathematical aspects of logic. Logic is a tool for critical reflection and analysis, but also for clear, precise capturing or representation of knowledge claims. In my research one of my primary goals – as I tried to illustrate above - has always been to show that incorporating also mathematical logic into the practising of cognitive philosophy, and thus going far beyond the usual so-called 'philosophical logic', may result in an impressive refinement of both the critical and formal aspects of philosophy.

Philosophy in this century offers valuable opportunities for interdisciplinary research and teaching, while cognitive philosophy, in particular, has the added potential to serve academic philosophers and philosophy students in meeting new challenges posed by the information age. Moreover, by including at least psychology, computer science, linguistics, neuroscience and perhaps mathematics, cognitive philosophy can be expanded to cognitive science studies.[17].

[17] In general, the cognitive science student can communicate clearly, critically evaluate evidence and data, program and use computers, and apply conceptual, analytical, and interpersonal skills to a variety of situations, guided by their ability to weigh moral decisions clearly and with accountability. All of these represent skills that are extremely important both at undergraduate and postgraduate level, as well as in the active persuasion of a career. Students are generally given training in the nature of knowledge and intelligence (philosophy of mind, epistemology, philosophy of artificial intelligence, linguistics, sociology, anthropology, communication), in the neural mechanics underlying complex capacities (neuroscience, cognitive psychology, philosophy of mind, epistemology, artificial intelligence), in the computational techniques under development to model these capacities (computer science, logic, philosophy of artificial intelligence), as well as in dealing with the moral implications implied by the fields covered by cognitive science studies (certain modules in applied philosophy). In this way students gain the conceptual framework and technical skills necessary for careers in research, teaching, information science, business, and government.

Currently, in SA, the University if KwaZulu-Natal is the only institution I know of that offers a specific qualification in cognitive science. They offer a BA and a BA(Honours) degree in cognitive science, with students drawing modules mainly from philosophy, linguistics, psychology, and computer science. Cognitive philosophers can play a leading role in satisfying the urgent demand for the establishment of postgraduate qualifications in cognitive science studies.

CONCLUSION

Many of you may still see philosophy as a 'woven web of guesses' (as Popper referred to science) without any practical application. The fact of the matter is though, that philosophy as the most fundamental discipline of them all, still, after more than 2000 years after Socrates's birth, offers us the best toolkit for investigating the origin, nature, products, and evolution of the processes of knowledge. Moreover, the fact that philosophy still plays, at least in terms of studies of mind and consciousness, its traditional role of clarifying the field before science steps in; combined with its many foundational roles in the sciences; its practical application today in terms of some of the more viable and respectable branches of 'applied philosophy'; and, most importantly, its application in the field of cognitive science, show that no-one need despair yet about its future.

Philosophy is about beginnings, not about endings. It is about uncertainty as much as it is about certainty, because it is about critique and revision in the light of new knowledge. Socrates turned uncertainty into a virtue, and so it should remain. I became interested in Philosophy as a teenager because the Greek idea of the beauty of form and precision attracted me. I now know that this notion of 'beauty and precision of form' is transient, but always worthwhile; rather than getting lost in the flux produced by philosophical criticism, it remains the goal of Socratic uncertainty and critique.

In Hilary Putnam's (2000) last lecture at Harvard before he retired in 2000, he said: '… philosophy is a good thing. It's a great thing. It can lead to wonderful things, and it can lead to terrible things. But it means … that you take the responsibility of trying to think deeply and with integrity seriously'. Thank you.

REFERENCES

Carnap, R. 1966. *Philosophical foundations of physics: An introduction to the philosophy of science.* New York: Basic Books, Inc..

Carnap, R. 1975. Observation language and theoretical language, in: HINTIKKA, J. (Ed.) *Rudolf Carnap, logical empiricist*. Dordrecht: D. Reidel Publishing Company, 75-86.

Cartwright, N. 1983. *How the laws of physics lie.* Oxford: Oxford University Press.

Delgrande, J., T. Schaub, H. Tompits, & K. Wang 2004. A classification and survey of preference handling approaches in nonmonotonic reasoning, *Computational Intelligence.* 20(2),308-334.

Holton, G. 1995. The role of themata in science, *Foundations in physics.* 26(4), 453-465.

Kuipers, T.A.F. 2000. *From instrumentalism to constructive realism*. Dordrecht: Kluwer Academic Publishers.

Labuschagne, W.A. 1994. *Truth*. Inaugural Lecture, Department of Computer Science, University of South Africa. 1994.

Lötter, H.P.P. 1994. A postmodern philosophy of science?, *South African journal of philosophy*. 13(3), 153-160.

Makinson, D. 2005. *Bridges from classical to non-monotonic logic. Texts in computing vol.5*. London: King's College.

Nagel, E. 1949. The meaning of reduction in the natural sciences, in: R.C. Stauffer (Ed.) *Science and civilisation.* Madison: University of Wisconsin Press.

Nagel, E. 1961. *The structure of science.* London: Routledge & Kegan Paul.

Niiniluoto, I. 1999. *Critical scientific realism.* Oxford: Oxford University Press.

Putnam, H. 1983. Models and reality, in: H. Putnam *Realism and reason.* New York: Cambridge University Press, 1-25.

Putnam, H. 2000. 'To think with integrity', *The Harvard Review of Philosophy.* VIII, 4-13.

Ruttkamp, E.B. 2002. *A model-theoretic realist interpretation of science.* Dordrecht: Kluwer Academic Publishers.

Ruttkamp, E.B. 2005. 'Over determination of theories by empirical models: A realist interpretation of empirical choices', In: Festa, R., A. Aliseda & J. Peijnenburg (Eds.) *Cognitive structures in scientific inquiry.* Amsterdam: Rodopi, Poznań Studies in the Philosophy of the Sciences and Humanities, 84, 409-436.

Ruttkamp, E.B. 2006. 'Reduction revisited', *South African Journal of Philosophy*, 25(2), 102-112.

Ruttkamp, E.B. & J. Heidema. 2005. Reviewing reduction in a preferential model-theoretic context, *International Studies in Philosophy of Science.* 19(2), 123-146.

Shapiro, S. 2000. *Classical logic.* The Stanford Encyclopaedia of Philosophy (Fall 2005 Edition), Edward N. Zolta (Ed.), URL = http://plato.stanford.edu/archives/fall2005/entries/logic-classical.

Shoham, Y. 1988. *Reasoning about change: time and causation from the standpoint of artificial intelligence.* Cambridge MA: MIT Press.

Sowa, J.F. 2000. *Knowledge Representation. Logical, philosophical, and computational foundations.* London: Brooks/Cole Thompson Learning.

Tarski, A. 1935. Der Wahrheitsbegriff in formalisierten Sprachen, translated as The concept of truth in formalized languages, in Tarksi, A. 1982. (Second Edition) *Logic, semantics, metamathematics.* Indianapolis: Hackett Publishing Co, 152-278.

Tarksi, A. 1982. (Second Edition) *Logic, semantics metamathematics.* Indianapolis: Hackett Publishing Co.

Suppes, P. 2002. *Representation and invariance of scientific structures.* Stanford: CSLI Publications.

Van Fraassen, B.C. 1980. *The scientific image*. Oxford: Oxford University Press.