

Models and Representation

Roman Frigg¹ and James Nguyen²

Forthcoming in Lorenzo Magnani and Tommaso Bertolotti (eds.): Springer Handbook of Model-Based Science, Berlin and New York: Springer 2017.³

1. Introduction

Models play a central role in contemporary science. Scientists construct models of atoms, elementary particles, polymers, populations, genetic trees, economies, rational decisions, aeroplanes, earthquakes, forest fires, irrigation systems, and the world's climate – there is hardly a domain of inquiry without models. Models are essential for the acquisition and organisation of scientific knowledge. We often study a model to discover features of the thing it stands for. How does this work? The answer is that a model can instruct us about the nature of its subject matter if it represents the selected part or aspect of the world that we investigate. So if we want to understand how models allow us to learn about the world, we have to come to understand how they represent.

The problem of representation has generated a sizable literature, which has been growing fast in particular over the last decade. The aim of this chapter is to review this body of work and assess the strengths and weaknesses of the different proposals. This enterprise faces an immediate difficulty: even a cursory look at the literature on scientific representation quickly reveals that there is no such thing as 'the' problem of scientific representation. In fact, we find a cluster of interrelated problems. In Section

¹ Contact r.p.frigg@lse.ac.uk, Department of Philosophy, Logic and Scientific Method and Centre for Philosophy of Natural and Social Science; London School of Economics and Political Science.

² Contact jnguyen4@nd.edu History and Philosophy of Science; University of Notre Dame and Centre for Philosophy of Natural and Social Science; London School of Economics and Political Science.

³ The published version will be available on www.romanfrigg.org as soon as the book has come out.

2 we try to untangle this web and get clear on what the problems are and on how they relate to one another (for a historical introduction to the issue see Boniolo (2007)). The result of this effort is a list with five problems and five conditions of adequacy, which provides the analytical lens through which we look at the different accounts. In Section 3 we discuss Griceanism and *stipulative fiat*. In Section 4 we look at the time-honoured similarity approach, and in Section 5 we examine its modern-day cousin, the structuralist approach. In Section 6 we turn to inferentialism, a more recent family of conceptions. In Section 7 we discuss the fiction view of models, and in Section 8 we consider the conception of representation-as.

Before delving into the discussion, a number of caveats are in order. The first is that our discussion in no way presupposes that models are the sole unit of scientific representation, or that all scientific representation is model-based. Various types of images have their place in science, and so do graphs, diagrams, and drawings (Perini (2005; 2005; 2010) and Elkins (1999) provide discussions of visual representation in the sciences). In some contexts scientists use what Warmbröd (1992) calls ‘natural forms of representation’ and what Peirce (1932) would have classified as indices: tree rings, fingerprints, disease symptoms. These are related to thermometer readings and litmus paper indications, which are commonly classified as measurements. Measurements also provide representations of processes in nature, sometimes together with the subsequent condensation of measurement results in the form of charts, curves, tables and the like (Tal (2015) provides a discussion of measurement). And, last but not least, many would hold that theories represent too. At this point the vexing problem of the nature of theories and the relation between theories and models rears its head again. We refer the reader to Portides’ contribution to this volume (forthcoming) for a discussion of this issue. Whether these other forms of scientific representation have features in common with how models represent is an interesting question, but this is a problem for another day. Our aim here is a more modest one: to understand how models represent. To make the scope of our investigation explicit we call the kind of representation we are interested in ‘model-representation’.

The second point to emphasise is that our discussion is not premised on the claim that *all* models are representational; nor does it assume that representation is the only (or even primary) function of models. It has been emphasised variously that models

perform a number of functions other than representation. To mention but few: Knuuttila (2005; 2011) points out that the epistemic value of models is not limited to their representational function and develops an account that views models as epistemic artifacts which allow us to gather knowledge in diverse ways; Morgan and Morrison (1999) emphasise the role models play in the mediation between theories and the world; Hartmann (1995) discusses models as tools for theory construction; Peschard (2011) investigates the way in which models may be used to construct other models and generate new target systems; and Bokulich (2009) and Kennedy (2012) present non-representational accounts of model explanation (Woody (2004) and Reiss (2012) provide general discussions of the relation between representation and explanation). Not only do we not see projects like these as being in conflict with a view that sees some models as representational; we think that the approaches are in fact complementary.

Finally, there is a popular myth according to which a representation is a mirror image, a copy, or an imitation of the thing it represents. On this view representation is *ipso facto* realistic representation. This is a mistake. Representations can be realistic, but they need not. And representations certainly need not be copies of the real thing. This, we take it, is the moral of the satire about the cartographers who produce maps as large as the country itself only to see them abandoned. The story has been told by Lewis Carroll in *Sylvie and Bruno* and Jorge Luis Borges in *On Exactitude in Science*. Throughout this review we encounter positions that make room for non-realistic representation and hence testify to the fact that representation is a much broader notion than mirroring.

There is, however, a sense in which we presuppose a minimal form of realism. Throughout the discussion we assume that target systems exist independently of human observers, and that they are how they are irrespective of what anybody thinks about them. That is, we assume that the targets of representation exist independently of the representation. This is a presupposition not everybody would share. Constructivists (and other kinds of metaphysical antirealists) assume that there is no phenomenon independent of its representation: representations constitute the phenomena they represent (this view is expounded for instance by Lynch and Wooglar (1990); Giere (1994) offers a critical discussion). It goes without saying that

an assessment of the constructivist programme is beyond the scope of this review. It is worth observing, though, that many of the discussions to follow are by no means pointless from a constructivist perspective. What in the realist idiom is conceptualised as the representation of an object in the world by a model would, from the constructivist perspective, turn into the study of the relation between a model and another representation, or an object constituted by another representation. This is because even from a constructivist perspective, models and their targets are not identical, and the fact that targets are representationally constituted would not obliterate the differences between a target representation and scientific model.

2. Problems Concerning Model-Representation

In this section we say what questions a philosophical account of model-representation has to answer and reflect on what conditions such an answer has to satisfy. As one would expect, different authors have framed the problem in different ways. Nevertheless, recent discussions about model-representation have tended to cluster around a relatively well-circumscribed set of issues. The aim of this section is to make these issues explicit and formulate five problems that an account of model-representation has to answer. These problems will help us structuring the discussion in later sections and put views and positions into perspective. In the course of doing so we also articulate five conditions of adequacy that every account of model-representation has to satisfy.

Models are representations of a selected part or aspect of the world. This is the model's *target system*. The first and most fundamental question about a model therefore is: in virtue of what is a model a representation of something else? Attention has been drawn to this issue by Frigg (2002, p. 17; 2006, p. 50), Morrison (2008, p. 70), and Suárez (2003, p. 230). To appreciate the thrust of this question it is instructive to briefly ponder the same problem in the context of pictorial representation. When seeing, say, Soutine's *The Groom or the Bellboy* we immediately realise that it depicts a man in a red dress. Why is this? *Per se* the painting is a plane surface covered with pigments. How does an arrangement of pigments on a surface represent something outside the picture frame? Likewise,

models, before being representations of atoms, populations, or economies, are equations, structures, fictional scenarios, or mannerly physical objects. The problem is: what turns equations and structures, or fictional scenarios and physical objects into representations of something beyond themselves? It has become customary to phrase this problem in terms of necessary and sufficient conditions and throughout this review we shall follow suit (some may balk at this, but it's worth flagging that the standard arguments against such an analysis, e.g. those surveyed in Laurence and Margolis's (1999), lose much of their bite when attention is restricted to core cases as we do here). The question then is: What fills the blank in ' M is a model-representation of T iff ___', where ' M ' stands for 'model' and ' T ' for 'target system'?

To spare ourselves difficulties further down the line, this formulation needs to be adjusted in light of a crucial condition of adequacy that any account of model-representation has to meet. The condition is that models represent in a way that allows us to form hypotheses about their target systems. We can generate claims about a target system by investigating a model that represents it. Many investigations are carried out on models rather than on reality itself, and this is done with the aim of discovering features of the things models stand for. Every acceptable theory of scientific representation has to account for how reasoning conducted on models can yield claims about their target systems. Let us call this the *Surrogate Reasoning Condition*.

The term 'surrogate reasoning' was introduced by Swoyer (1991, p. 449), and there seems to be widespread agreement on this point (although Callender and Cohen (2006), whose views are discussed in Section 3, provide a noteworthy exception). To mention just a few writers on the subject: Bailer-Jones (2003, p. 59) emphasises that models '*tell us* something about certain features of the world' (original emphasis). Boliskna (2013) and Contessa (2007) both call models 'epistemic representations'; Frigg (2003, p. 104; 2006, p. 51) sees the potential for learning as an essential explanandum for any theory of representation; Liu (2013, p. 93) emphasises that the main role for models in science and technology is epistemic; Morgan and Morrison regard models as 'investigative tools' (1999, p. 11); Suárez (2003, p. 229; 2004, p. 772) submits that models licence specific inferences about their targets; and Weisberg (2013, p. 150) observes that the 'model-world relation is the relationship in virtue of

which studying a model can tell us something about the nature of a target system'. This distinguishes models from lexicographical representations such as words. Studying the internal constitution of a model can provide information about the target. Not so with words. The properties of a word (consisting of so and so many letters and syllables, occupying this or that position in a dictionary, etc.) do not matter to its functioning as a word; and neither do the physical properties of the ink used to print words on a piece of paper. We can replace one word by another at will (which is what happens in translations from one language to another), and we can print words with other methods than ink on paper. This is possible because the properties of a word as an object do not matter to its semantic function.

This gives rise to a problem for the schema ' M is a model-representation of T iff ____'. The problem is that any account of representation that fills the blank in a way that satisfies the *Surrogative Reasoning Condition* will almost invariably end up covering other kinds of representations too. Geographical maps, graphs, diagrams, charts, drawings, pictures, and photographs often provide epistemic access to features of the items they represent, and hence are likely to fall under an account of representation that explains this sort of reasoning. This is a problem for an analysis of model-representation in terms of necessary and sufficient conditions because if something that is not *prima facie* a model (for instance a map or a photograph) satisfies the conditions of an account of model-representation, then one either has to conclude that the account fails because it does not provide necessary conditions, or that first impressions are wrong and other representations (such as maps or photographs) are in fact model-representations.

Neither of these options is appealing. To avoid this problem we follow a suggestion of Contessa's (2007) and broaden the scope of the investigation. Rather than analysing the relatively narrow category of model-representation, we analyse the broader category of *epistemic representation*. This category comprises model-representations, but it also includes other representations that allow for surrogative reasoning. The task then becomes to fill the blank in ' M is an epistemic representation of T iff ____'. For brevity we use ' $R(M, T)$ ' as a stand in for ' M is an epistemic representation of T ', and so the biconditional becomes ' $R(M, T)$ iff ____'. We call the general problem of figuring out in virtue of what something is an epistemic representation of something

else the *Epistemic Representation Problem* (*ER-Problem*, for short), and the above biconditional the *ER-Scheme*. So one can say that the ER-Problem is to fill the blank in the ER-Scheme. Frigg (2006, p. 50) calls this the ‘enigma of representation’ and in Suárez’s (2003, p. 230) terminology this amounts to identifying the ‘constituents’ of a representation (although he questions whether both necessary *and* sufficient conditions can be given; see Section 6 for further discussion on how his views fit into the ER-framework).

Analysing the larger category of epistemic representation and placing model-representations in that category can be seen as giving rise to a demarcation problem for scientific representations: how do scientific model-representations differ from other kinds of epistemic representations? We refer to this question as the *Representational Demarcation Problem*. Callender and Cohen (2006, p. 69) formulate this problem, but then voice scepticism about our ability to solve it (*ibid.*, p. 83). The representational demarcation problem has received little, if any, attention in the recent literature on scientific representation, which would suggest that other authors either share Callender and Cohen’s scepticism, or regard it as a non-issue to begin with. The latter seems to be implicit in approaches which discuss scientific representation alongside pictorial representation such as Elgin’s (2010), French’s (2003), Frigg’s (2006), Suárez’s (2004), and van Fraassen’s (2008). But a dismissal of the problem is in no way neutral stance. It amounts to no less than the admission that model-representations are not fundamentally different from other epistemic representations, or that we are unable to pin down what the distinguishing features are. Such a stance should be made explicit and, ideally, justified.

Two qualifications concerning the ER-Scheme need to be added. The first concerns its flexibility. Some might worry that posing the problem in this way prejudices what answers can be given. The worry comes in a number of variants. A first variant is that the scheme presupposes that representation is an intrinsic relation between M and T (i.e. a relation that only depends on intrinsic properties of M and T and on how they relate to one another rather than on how they relate to other objects) or even that it is naturalisable (a notion further discussed in Section 4). This is not so. In fact, R might depend on any number of factors other than M and T themselves, and on ones that do not qualify as natural ones. To make this explicit we write the ER-Scheme in the form

' $R(M, T)$ iff $C(M, T, x_1, \dots, x_n)$ ', where n is a natural number and C is an $(n+2)$ -ary relation that grounds representation. The x_i can be anything that is deemed relevant to epistemic representation, for instance a user's intentions, standards of accuracy, and specific purposes. We call C the *grounding relation* of an epistemic representation.

Before adding a second qualification, let us introduce the next problem in connection with model-representation. Even if we restrict our attention to scientific epistemic representations (if they are found to be relevantly different to non-scientific epistemic representations as per the demarcation problem above), not all representations are of the same kind. In the case of visual representations this is so obvious that it hardly needs mention: an Egyptian mural, a two-point perspective ink drawing, a pointillist oil painting, an architectural plan, and a road map represent their respective targets in different ways. This pluralism is not limited to visual representations. Model-representations do not all seem to be of the same kind either. Woody (2000) argues that chemistry as a discipline has its own ways to represent molecules. But differences in style can also appear in models from the same discipline. Weizsäcker's liquid drop model represents the nucleus of an atom in a manner that seems to be different from the one of the shell model. A scale model of the wing of a plane represents the wing in a way that is different from how a mathematical model of its cross section does. Or Phillips and Newlyn's famous hydraulic machine and Hicks' mathematical models both represent a Keynesian economy but they seem to do so in different ways. This gives rise to the question: what styles are there and how can they be characterised? This is the *Problem of Style* (Frigg 2006, p. 50). There is no expectation that a *complete* list of styles be provided in response. Indeed, it is unlikely that such a list can ever be drawn up, and new styles will be invented as science progresses. For this reason a response to the problem of style will always be open-ended, providing a taxonomy of what is currently available while leaving room for later additions.

With this in mind we can now turn to the second qualification concerning the ER-Scheme. The worry is this: the scheme seems to assume that representation is a monolithic concept and thereby make it impossible to distinguish between different kinds of representation. The impression is engendered by the fact the scheme asks us

to fill a blank, and blank is filled only once. But if there are different kinds of representations, we should be able to fill the blank in different ways on different occasions because a theory of representation should not force upon us the view that the different styles are all variations of one overarching concept of representation.

The ER-Scheme is more flexible than it appears at first sight. There are at least three ways in which different styles of representations can be accommodated. For the sake of illustration, and to add some palpability to an abstract discussion, let us assume that we have identified two styles: analogue representation and idealised representation. The result of an analysis of these relations is the identification of their respective grounding relations. Let $C_A(M, T, \dots)$ and $C_I(M, T, \dots)$ be these relations. The first way of accommodating them in the ER-Scheme is to fill the blank with the disjunction of the two: ' $R(M, T)$ iff $C_A(M, T, \dots)$ or $C_I(M, T, \dots)$ '. In plain English: M represents T if and only if M is an analogue representation of T or M is an idealised representation of T . This move is possible because, first appearances notwithstanding, nothing hangs on the grounding relation being homogenous. The relation can be as complicated as we like and there is no prohibition against disjunctions. In the above case we have $C = [C_A \text{ or } C_I]$. Furthermore, the grounding relation could even be an open disjunction. This would help accommodating the above observation that a list of styles is potentially open-ended. In that case there would be a grounding relation for each style and the scheme could be written as ' $R(M, T)$ iff $C_1(M, T \dots)$ or $C_2(M, T \dots)$ or $C_3(M, T \dots)$ or ...', where the C_i are the grounding relations for different styles. This is not a new scheme; it's the old scheme where $C = [C_1 \text{ or } C_2 \text{ or } C_3 \text{ or } \dots]$ is spelled out.

Alternatively one could formulate a different scheme for every kind of representation. This would amount to changing the scheme slightly in that one does not analyse epistemic representation *per se*. Instead one would analyse different kinds of epistemic representations. Consider the above example again. Let $R_1(M, T)$ stand for ' M is an analogue epistemic representation of T ' and $R_2(M, T)$ for ' M is an idealised epistemic representation of T '. The response to the ER-Problem then consists in presenting the two biconditionals ' $R_1(M, T)$ iff C_A ' and ' $R_2(M, T)$ iff C_I '. This generalises straightforwardly to the case of any number of styles, and the open-endedness of the list of styles can be reflected in the fact that an open-ended list of

conditionals of the form ‘ $R_i(M, T)$ iff C_i ’ can be given, where the index ranges over styles.

In contrast with the second option, which pulls in the direction of more diversity, the third aims for more unity. The crucial observation here is that the grounding relation can in principle be an abstract relation that can be concretised in different ways, or a determinable that can have different determinates. On the third view, then, the concept of representation is like the concept of force (which is abstract in that in a concrete situation force is gravity or electromagnetic attraction or some other specific force), or like colour (where a coloured object must be blue or green or ...). This view would leave ‘ $R(M, T)$ iff $C(M, T, x_1, \dots, x_n)$ ’ unchanged and take it as understood that C is an abstract relation.

At this point we do not adjudicate between these options. Each has its own pros and cons, and which one is the most convenient to work with depends on one’s other philosophical commitments. What matters is that the ER-scheme does have the flexibility to accommodate different representational styles, and that it can in fact accommodate them in at least three different ways.

The next problem in line for theory of model-representation is to specify standards of accuracy. Some representations are accurate; others aren’t. The Schrödinger model is an accurate representation of the hydrogen atom; the Thomson model isn’t. On what grounds do we make such judgments? In Morrison’s words: ‘how do we identify what constitutes an accurate representation?’ (2008, p. 70). We call this the Problem of *Standards of Accuracy*. Answering this question might make reference to the purposes of the model and model user, and thus it is important to note that by ‘accuracy’ we mean something that can come in degrees and may be context dependant. Providing a response to the problem of accuracy is a crucial aspect of an account of epistemic representation.

This problem goes hand in hand with a second condition of adequacy: the *Possibility of Misrepresentation*. Asking what makes a representation an accurate representation already presupposes that inaccurate representations are representations too. And this

is how it should be. If M does not accurately portray T , then it is a misrepresentation but not a non-representation. It is therefore a general constraint on a theory of epistemic representation that it has to make misrepresentation possible. This can be motivated by a brief glance at the history of science, but is plausibly also part of the concept of representation, and as such is found in discussions of other kinds of representation (Stitch and Warfield (1994, pp. 6-7), for instance, suggest that a theory of mental representation should be able to account for misrepresentation, as do Sterelny and Griffiths (1999, p. 104) in their discussion of genetic representation). A corollary of this requirement is that representation is a wider concept than accurate representation and that representation cannot be analysed in terms of accurate representation.

A related condition concerns models that misrepresent in the sense that they lack target systems. Models of ether, phlogiston, four-sex populations, and so on, are all deemed scientific models, but ether, phlogiston, and four-sex populations don't exist. Such models lack (actual) target systems, and one hopes that an account of epistemic representation would allow us to understand how these models work. We call this the problem of targetless models (or models without targets).

The fourth condition of adequacy for an account of model-representation is that it must account for the directionality of representation. Models are about their targets, but (at least in general) targets are not about their models. So there is an essential directionality to representations, and an account of model-representation has to identify the root of this directionality. We call this the *Requirement of Directionality*.

Many scientific models are highly mathematized, and their mathematical aspects are crucial to their cognitive as well as their representational function. This forces us to reconsider a time-honoured philosophical puzzle: the applicability of mathematics in the empirical sciences. Even though the problem can be traced back at least to Plato's *Timaeus*, its canonical modern expression is due to Wigner, who famously remarked that 'the enormous usefulness of mathematics in the natural sciences is something bordering on the mysterious and that there is no explanation for it' (1960, p. 2). One need not go as far as seeing the applicability of mathematics as an inexplicable miracle, but the question remains: how does mathematics hook onto the world?

The recent discussion of this problem has taken place in a body of literature that grew out of the philosophy of mathematics (see Shapiro (1997, Chapter 8) for a review). But, with the exception of Bueno and Colyvan's (2011), there has been little contact with the literature on scientific modelling. This is a regrettable state of affairs. The question how a mathematized model represents its target implies the question how mathematics applies to a physical system. So rather than separating the question of model-representation from the problem of the applicability of mathematics and dealing with them in separate discussions, they should be seen as the two sides of the same coin and be dealt with in tandem. For this reason, our fifth and final condition of adequacy is that an account of representation has to explain how mathematics is applied to the physical world. We call this the *Applicability of Mathematics Condition*.

In answering the above questions one invariably runs up against a further problem, the *Problem of Ontology*: what kinds of objects are models? Are they structures in the sense of set theory, fictional entities, descriptions, equations or yet something else? Or are there no models at all? While some authors develop an ontology of models, others reject an understanding of models as 'things' and push a programme that can be summed up in the slogan 'modelling without models' (Levy 2015). There is also no presupposition that all models be of the same kind. Some models are material objects, some are things that one holds in one's head rather than one's hands (to use Hacking's phrase (1983, p. 216)). For the most part, the focus in debates about representation has been on non-material models, and we will follow this convention. It is worth emphasising, however, that also the seemingly straightforward material models raise interesting philosophical questions: Rosenblueth and Wiener (1945) discuss the criteria for choosing an object as a model; Ankeny and Leonelli (2011) discuss issues that arise when using organisms as models; and the contributors to (Klein 2001) discuss representation in the laboratory.

A theory of representation can recognise different kinds of models, or indeed no models at all. The requirement only asks us to be clear on our commitments and provide a list with things, if any, that we recognise as models and give an account of

what they are in case these entities raise questions (what exactly do we mean by something that one holds in one's head rather than one's hands?).

In sum, an account of model-representation has to do the following:

- (1) Provide an answer to the *Epistemic Representation Problem* (filling the blank in ER-Scheme: ' M is an epistemic representation of T iff ____').
- (2) Take a stand on the *Representational Demarcation Problem* (the question of how scientific epistemic representations differ from other kinds of epistemic representations).
- (3) Respond to the *Problem of Style* (what styles are there and how can they be characterised?).
- (4) Formulate *Standards of Accuracy* (how do we identify what constitutes an accurate representation?).
- (5) Address the *Problem of Ontology* (what kinds of objects are models?).

Any satisfactory answer to these five issues will have to meet the following five conditions of adequacy:

- (1) *Surrogate Reasoning Condition* (models represent their targets in a way that allows us to generate hypotheses about them).
- (2) *Possibility of Misrepresentation* (if M does not accurately represent T , then it is a misrepresentation but not a non-representation).
- (3) *Targetless Models* (what are we to make of scientific representations that lack targets?).
- (4) *Requirement of Directionality* (models are about their targets, but targets are not about their models).
- (5) *Applicability of Mathematics Condition* (how does the mathematical apparatus used in M latch onto the physical world).

To frame the problem in this way is not to say that these are separate and unrelated issues, which can be dealt with one after the other in roughly the same way in which we first buy a ticket, walk to the platform and then take a train. This division is analytical, not factual. It serves to structure the discussion and to assess proposals; it

does not imply that an answer to one of these questions can be dissociated from what stance we take on the other issues.

3. General Griceanism and Stipulative Fiat

Callender and Cohen (Callender and Cohen 2006) submit that the entire debate over scientific representation has started on the wrong foot. They claim that scientific representation is not different from ‘artistic, linguistic, and culinary representation’ and in fact ‘there is no special problem about scientific representation’ (*ibid.* p. 67). Underlying this claim is a position Callender and Cohen call ‘General Griceanism’ (GG). The core of GG is the reductive claim that most representations we encounter are ‘derivative from the representational status of a privileged core of representations’ (*ibid.*, p. 70). GG then comes with a practical prescription about how to proceed with the analysis of a representation: ‘the General Gricean view consists of two stages. First, it explains the representational powers of derivative representations in terms of those of fundamental representations; second, it offers some other story to explain representation for the fundamental bearers of content.’ (*ibid.*, p. 73) Of these stages only the second requires serious philosophical work, and this work is done in the philosophy of mind because the fundamental form of representation is mental representation.

Scientific representation is a derivative kind of representation (*ibid.*, p. 71,75) and hence falls under the first stage of the above recipe. It is reduced to mental representation by an act of stipulation:

‘Can the salt shaker on the dinner table represent Madagascar? Of course it can, so long as you stipulate that the former represents the latter. [...] Can your left hand represent the Platonic form of beauty? Of course, so long as you stipulate that the former represents the latter. [...] On the story we are telling, then, virtually anything can be stipulated to be a representational vehicle for the representation of virtually anything [...]; the representational powers of mental states are so wide-ranging that they can bring about other representational relations between arbitrary relata by dint of mere stipulation. The upshot is that, once one has paid the admittedly hefty one-time fee of supplying a metaphysics of representation for mental states, further instances of representation become extremely cheap.’ (*ibid.*, pp. 73-74)

So explaining any form of representation other than mental representation is a triviality – all it takes is an act of ‘stipulative fiat’ (*ibid.*, p. 75). This supplies their answer to the ER-problem:

Stipulative Fiat: A scientific model M represents a target system T iff a model user stipulates that M represents T .

On this view, scientific representations are cheap to come by. The question therefore arises why scientists spend a lot of time constructing and studying complex models if they might just as well take a salt shaker and turn it into a representation of, say, a Bose-Einstein condensate by an act of fiat. Callender and Cohen admit that there are useful and not so useful representations, and that salt shakers belong the latter group. However, they insist that this has nothing to do with representation: ‘the questions about the utility of these representational vehicles are questions about the pragmatics of things that are representational vehicles, not questions about their representational status *per se.*’ (*ibid.*, p. 75) So, in sum, scientific representation ‘is constituted in terms of a stipulation, together with an underlying theory of representation for mental states, isomorphism, similarity, and inference generation are all idle wheels’ (*ibid.*, p. 78).

The first question we are faced with when assessing this account is the relation between GG and Stipulative Fiat. Callender and Cohen do not comment on this issue, but that they mention both in the same breath would suggest that they regard them as one and the same doctrine, or at least as the two sides of the same coin. This is not so. Stipulative Fiat is just one way of fleshing out GG, which only requires that there be *some* explanation of how derivative representations relate to fundamental representations; GG does not require that this explanation be of a particular kind, much less that it consist in nothing but an act of stipulation (Toon 2010, pp. 77-78; 2012, p. 244). Even if GG is correct, it doesn’t follow that Stipulative Fiat is a satisfactory answer to the ER-problem. Model-representation can, in principle, be ‘reduced’ to fundamental representation in many different ways (some of which we will encounter later in this chapter). Conversely, the failure of Stipulate Fiat does not entail that we must reject GG: one can uphold the idea that an appeal to the intentions

of model users is a crucial element in an account of scientific representation even if one dismisses Stipulative Fiat.

Let us now examine Stipulative Fiat. Callender and Cohen emphasise that anything can be a representation of anything else (*ibid.*, p. 73). This is correct. Things that function as models don't belong to a distinctive ontological category, and it would be a mistake to think that that some objects are, intrinsically, representations and other are not. This point has been made by others too (including Frigg (2010, p. 99), Giere (2010, p. 269), Suárez (2004, p. 773), Swoyer (1991, p. 452), and Teller (2001, p. 397)) and, as we shall see, it is a cornerstone of several alternative accounts of representation.

But just because anything can, in principle, be a representation of anything else, it doesn't follow that a mere act of stipulation suffices to turn M into a representation of T . Furthermore, it doesn't follow that an object elevated to the status of a representation by an act of fiat represents its target in a way that can appropriately be characterised as an instance of epistemic representation. We discuss both concerns in reverse order.

Stipulative Fiat fails to meet the Surrogate Reasoning Condition: it fails to provide an account of how claims about Madagascar could be extracted from reasoning about the salt shaker. Even if we admit that Stipulative Fiat establishes that models denote their targets (and as we will see soon, there is a question about this), denotation is not sufficient for epistemic representation. Both the word 'Napoleon' and Jacques-Louis David's portrait of Napoleon serve to denote the French general. But this does not imply that they represent him in the same way, as noted by Toon (2010, pp. 78-79). Bueno and French (2011, pp. 871-874) gesture in the same direction when they point to Peirce's distinction between icon, index and symbol and dismiss Callender and Cohen's views on grounds that they cannot explain the obvious differences between different kinds of representations.

Supporters of Stipulative Fiat could try to mitigate the force of this objection in two ways. First, they could appeal to additional facts about the object, as well as its relation to other items, in order to account for surrogate reasoning. For instance, the

salt shaker being to the right of the pepper mill might allow us to infer that Madagascar is to the east of Mozambique. Moves of this sort, however, invoke (at least tacitly) a specifiable relation between features of the model and features of the target (similarity, isomorphism, or otherwise), and an invocation of this kind goes beyond mere stipulation. Second, the last quotation from Callender and Cohen suggests that they might want to relegate surrogative reasoning into the realm of pragmatics and deny that it is part of the relation properly called epistemic representation. This, however, in effect amounts to a removal of the Surrogative Reasoning Condition from the desiderata of an account of scientific representation, and we have argued in Section 2 that surrogative reasoning is one of the hallmarks of scientific representation. And even if it were ‘pragmatics’, we still would want an account of how it works.

Let us now turn to our first point, that a mere act of stipulation is insufficient to turn M into a representation of T . We take our cue from a parallel discussion in the philosophy of language, where it has been pointed out that it is not clear that stipulation is sufficient to establish a denotational relationship (which is weaker than epistemic representation). A position similar to Stipulative Fiat faces what is known as the *Humpty Dumpty problem*, named in reference to Lewis Carroll’s discussion of Humpty using the word ‘glory’ to mean ‘a nice knockdown argument’ (MacKay 1968) *cf.* (Donnellan 1968) (it’s worth noting that this debate concerns meaning, rather than denotation, but it’s plausible that it can be reconstructed in terms of the latter). If stipulation is all that matters, then as long as Humpty simply stipulates that ‘glory’ means ‘a nice knockdown argument’, then it does so. And this doesn’t seem to be the case. Even if the utterance ‘glory’ could mean ‘a nice knockdown argument’ – if, for example, Humpty was speaking a different language – in the case in question it doesn’t, irrespective of Humpty’s stipulation. In the contemporary philosophy of language the discussion of this problem focuses more on the denotation of demonstratives rather than proper names, and work in that field focuses on propping up existing accounts so as to ensure that a speaker’s intentions successfully establish the denotation of demonstratives uttered by the speaker (Michaelson 2013). Whatever the success of these endeavours, their mere existence shows that successfully establishing denotation requires moving beyond a bare appeal to stipulation, or brute intention. But if a brute appeal to intentions fails in the case of demonstratives – the

sorts of terms that such an account would most readily be applicable to – then we find it difficult to see how Stipulative Fiat will establish a representational relationship between models and their targets. Moreover, this whole discussion supposed that an intention-based account of denotation is the correct one. This is controversial – see Reimer and Michaelson’s (2014) for an overview of discussions of denotation in the philosophy of language. If this is not the correct way to think about denotation, then Stipulative Fiat will fail to get off the ground at all.

It now pays that we have separated GG from Stipulative Fiat. Even though Stipulative Fiat does not provide an adequate answer to the ER-problem, one can still uphold GG. As Callender and Cohen note, all that it requires is that there is a privileged class of representations (they take them to be mental states but are open to the suggestion that they might be something else (*ibid.*, p. 82)), and that other types of representations owe their representational capacities to their relationship with the primitive ones. So philosophers need an account of how members of this privileged class of representations represent, and how derivative representations, which includes scientific models, relate to this class.

This is a plausible position, and when stated like this, many recent contributors to the debate on scientific representation can be seen as falling under the umbrella of GG. As we will see below, the more developed versions of the similarity (Section 4) and isomorphism (Section 5) accounts of scientific representation make explicit reference to the intentions and purposes of model users, even if their earlier iterations did not. And so do the accounts discussed in the latter sections, where the intentions of model users (in a more complicated manner than that suggested by Stipulative Fiat) are invoked to establish epistemic representation.

4. The Similarity Conception

Moving on from the Gricean account we now turn to the similarity conception of scientific representation (in aesthetics the term ‘resemblance’ is used more commonly than ‘similarity’, but there does not seem to be a substantive difference between the notions, and we use the terms as synonyms throughout). Similarity and representation

initially appear to be two closely related concepts, and invoking the former to ground the latter has a philosophical lineage stretching back at least as far as Plato's *The Republic*.

In its most basic guise the similarity conception of scientific representation asserts that scientific models represent their targets in virtue of being similar to them. This conception has universal aspirations in that it is taken to account for epistemic representation across a broad range of different domains. Paintings, statues, and drawings are said to represent by being similar to their subjects, (see Abell's (2009) and Lopes's (2004) for relatively current discussions of similarity in the context of visual representation). And recently Giere, one of the view's leading contemporary proponents, proclaimed that it covers scientific models alongside 'words, equations, diagrams, graphs, photographs, and, increasingly, computer-generated images' (Giere 2004, p. 243) (see also Giere's (1996, p. 272), and for further discussion Toon's (2012, pp. 249-250)). So the similarity view repudiates the demarcation problem and submits that the same mechanism, namely similarity, underpins different kinds of representation in a broad variety of contexts. (Sometimes the similarity view is introduced by categorising models as icons in Peirce's sense, and, as Kraleman and Lattmann point out, icons represent 'on the basis of a similarity relation between themselves and their objects' (Kraleman and Lattmann 2013, p. 3398).)

The view also offers an elegant account of surrogative reasoning. Similarities between model and target can be exploited to carry over insights gained in the model to the target. If the similarity between M and T is based on shared properties, then a property found in M would also have to be present in T ; and if the similarity holds between properties themselves, then T would have to instantiate properties similar to M (however, it is worth noting that this kind of knowledge transfer can cause difficulties in some contexts, Frigg *et al.* (2014) discuss these difficulties in the context of non-linear dynamic modeling).

However, appeal to similarity in the context of representation leaves open whether similarity is offered as an answer to the ER-Problem, the Problem of Style, or whether it is meant to set Standards of Accuracy. Proponents of the similarity account typically have offered little guidance on this issue. So we examine each option in turn

and ask whether similarity offers a viable answer. We then turn to the question of how the similarity view deals with the Problem of Ontology.

4.1 Similarity and ER-Problem

Understood as response to the ER-Problem, a similarity view of representation amounts to the following:

Similarity I: A scientific model M represents a target T iff M and T are similar.

A well-known objection to this account is that similarity has the wrong logical properties. Goodman (1976, p. 4-5) submits that similarity is symmetric and reflexive yet representation isn't. If object A is similar to object B , then B is similar to A . But if A represents B , then B need not (and in fact in most cases does not) represent A : the Newtonian model represents the solar system, but the solar system does not represent the Newtonian model. And everything is similar to itself, but most things do not represent themselves. So this account does not meet our third condition of adequacy for an account of scientific representation insofar as it does not provide a direction to representation. (Similar problems also arise in connection with other logical properties, e.g. transitivity; see Frigg's (2003, p. 31) and Suárez's (2003, pp. 232-233)).

Yaghmaie (2012) argues that this conclusion – along with the third condition itself – is wrong: epistemic representation is symmetric and reflexive (he discusses this in the context of the isomorphism view of representation, which we turn to in the next section, but the point applies here as well). His examples are drawn from mathematical physics, and he presents a detailed case study of a symmetric representation relation between quantum field theory and statistical mechanics. His case raises interesting questions, but even if one grants that Yaghmaie has identified a case where representation is reflexive and symmetrical it does not follow that representation *in general* is. The photograph in Jane's passport represents Jane; but Jane does not represent her passport photograph; and the same holds true for myriads of other representations. Goodman is correct in pointing out that typically

representation is not symmetrical and reflexive: a target T does not represent model M just because M represents T .

A reply diametrically opposed to Yaghmaie's emerges from the writings of Tversky and Weisberg. They accept that representation is not symmetric, but dispute that similarity fails on this count. Using a gradual notion of similarity (i.e. one that allows for statements like ' A is similar to B to degree d '), Tversky found that subjects in empirical studies judged that North Korea was more similar to China than China was to North Korea (Tversky and Gati 1978); similarly Poznic (2015, Sec. 4.2) points out with reference to the characters in a Polanski movie that the similarity relation between a baby and the father need not be symmetric.

So allowing degrees into ones notion of similarity makes room for an asymmetry (although degrees by themselves are not sufficient for asymmetry; metric-based notions are still symmetric). This raises the question of how to analyse similarity. We discuss this thorny issue in some detail in the next subsection. For now we concede the point and grant that similarity need not always be symmetrical. However, this does not solve Goodman's problem with reflexivity (as we will see on Weisberg's notion of similarity everything is maximally similar to itself); nor does it, as will see now, solve other problems of the similarity account.

However the issue of logical properties is resolved, there is another serious problem: similarity is too inclusive a concept to account for representation. In many cases neither one of a pair of similar objects represents the other. Two copies of the same book are similar but neither represents the other. Similarity between two items is not enough to establish the requisite relationship of representation; there are many cases of similarity where no representation is involved. And this won't go away even if similarity turns out to be non-symmetric. That North Korea is similar to China (to some degree) does not imply that North Korea represents China, and that China is not similar to North Korea to the same degree does not alter this conclusion.

This point has been brought home in now-classical thought experiment due to Putnam (Putnam 1981, pp. 1-3) (but see also Black's (1973, p. 104)). An ant is crawling on a patch of sand and leaves a trace that happens to resemble Winston Churchill. Has the

ant produced a picture of Churchill? Putnam's answer is that it didn't because the ant has never seen Churchill and had no intention to produce an image of him. Although *someone else* might see the trace as a depiction of Churchill, the trace itself does not represent Churchill. This, Putnam concludes, shows that '[s]imilarity [...] to the features of Winston Churchill is not sufficient to make something represent or refer to Churchill.' (*ibid.*, p. 1). And what is true of the trace and Churchill is true of every other pair of similar items: similarity on its own does not establish representation.

There is also a more general issue concerning similarity: it is too easy to come by. Without constraints on what counts as similar, any two things can be considered similar to any degree (Aronson et al. 1995, p. 21). This, however, has the unfortunate consequence that anything represents anything else because any two objects are similar in some respect. Similarity is just too inclusive to account for representation. An obvious response to this problem is to delineate a set of relevant respects and degrees to which M and T have to be similar. This suggestion has been made explicitly by Giere (1988, p. 81) who suggests that models come equipped with what he calls 'theoretical hypotheses', statements asserting that model and target are similar in relevant respects and to certain degrees. This idea can be moulded into the following definition:

Similarity 2: A scientific model M represents a target T iff M and T are similar in relevant respects and to the relevant degrees.

On this definition one is free to choose one's respects and degrees so that unwanted similarities drop out of the picture. While this solves the last problem, it leaves the others untouched: similarity in relevant respects and to the relevant degrees is reflexive (and symmetrical, depending on one's notion of similarity); and presumably the ant's trace in the sand is still similar to Churchill in the relevant respects and degrees but without representing Churchill. Moreover, *Similarity 2* introduces three new problems.

First, a misrepresentation is one that portrays its target as having properties that are not similar in the relevant respects and to the relevant degrees to the true properties of the target. But then, on *Similarity 2*, M is not a representation at all. Ducheyne (2008)

embraces this conclusion when he offers a variant of a similarity account which explicitly takes the *success* of the hypothesized similarity between a model and its target to be a necessary condition on the model representing the target. In Section 2 we argued that the possibility of misrepresentation is condition of adequacy for any acceptable account of representation and so we submit that misrepresentation should not be conflated with non-representation (*cf.* Frigg 2002, p. 16; Suárez 2003, p. 235).

Second, similarity in relevant respects and to the relevant degrees does not guarantee that M represents the right target. As Suárez points out (2003, p. 233-4), even a regimented similarity can obtain with no corresponding representation. If John dresses up as Pope Innocent X (and he does so perfectly), then he resembles Velázquez's portrait of the pope (at least in as far as the pope himself resembled the portrait). In cases like these, which Suárez calls 'mistargeting', a model represents one target rather than another, despite the fact that both targets are relevantly similar to the model. Like in the case of Putnam's ant, the root cause of the problem is that the similarity is accidental. In the case of the ant, the accident occurs at the representation-end of the relation, whereas in the case of John's dressing up the accidental similarity occurs at the target-end. Both cases demonstrate that *Similarity 2* cannot rule out accidental representation.

Third, there may simply be nothing to be similar to because some representations represent no actual object (Goodman 1976, p. 26). Some paintings represent elves and dragons, and some models represent phlogiston and the ether. None of these exist. As Toon points out, this is a problem in particular for the similarity view (Toon 2012, pp. 246-247): models without objects cannot represent what they seem to represent because in order for two things to be similar to each other both have to exist. If there is no ether, then an ether model cannot be similar to the ether.

It would seem that at least the second problem could be solved by adding the requirement that M denote T (as considered, but not endorsed, by Goodman (1976, pp. 5-6)). Amending the previous definition accordingly yields:

Similarity 3: A scientific model M represents a target T iff M and T are similar in relevant respects and to the relevant degrees and M denotes T .

This account would also solve the problem with reflexivity (and symmetry), because denotation is directional in way similarity is not. Unfortunately *Similarity 3* still suffers from the first and the third problems. It would still lead to the conflation of misrepresentations with nonrepresentations because the first conjunct (similar in the relevant respects) would still be false. And a non-existent system cannot be denoted and so we have to conclude that models of, say, the ether and phlogiston represent nothing. This seems an unfortunate consequence because there is a clear sense in which models without targets are about something. Maxwell's writings on the ether provide a detailed and intelligible account of a number of properties of the ether, and these properties are highlighted in the model. If ether existed then *Similarity 3* could explain why these were important by appealing to them as being relevant for the similarity between an ether model and its target. But since ether does not, no such explanation is offered.

A different version of the similarity view sets aside the moves made in *Similarity 3* and tries to improve on *Similarity 2*. The crucial move is to take the very act of *asserting* a specific similarity between a model and a target as constitutive of the scientific representation.

Similarity 4: A scientific model M represents a target system T if and only if a theoretical hypotheses H asserts that M and T are similar in certain respects and to certain degrees.

This comes close to the view Giere advocated in *Explaining Science* (Giere 1988, p. 81) (something like this is also found in Cartwright (1999, pp. 192-193; cf. 1999, pp. 261-262) who appeals to a 'loose notion of resemblance'; her account of modelling is discussed in more detail in subsection 7.3). This version of the similarity view avoids problems with misrepresentation because, being hypotheses, there is no expectation that the assertions made in H are true. If they are, then the representation is accurate (or the representation is accurate to the extent that they hold). If they are not, then the representation is a misrepresentation. It resolves the problem of mistargeting because hypotheses identify targets before asserting similarities with M (that is, the task of

picking the right target is now placed in the court of the hypothesis and is no longer expected to be determined by the similarity relation). Finally it also resolves the issue with directionality because H can be understood as introducing a directionality that is not present in the similarity relation. However, it fails to resolve the problem with representation without a target. If there is no ether, no hypotheses can be asserted about it.

Let us set the issue of non-existent targets aside for the moment and have a closer look at the notion of representation proposed in *Similarity 4*. A crucial point remains understated in *Similarity 4*. Hypotheses don't assert themselves; hypotheses are put forward by those who work with representations, in the case of models, scientists. So the crucial ingredient – users – is left implicit in *Similarity 4*.

In a string of recent publications Giere made explicit the fact that 'scientists are intentional agents with goals and purposes' (Giere 2004, p. 743) and proposed to build this insight explicitly into an account of epistemic representation. This involves adopting an agent-based notion of representation that focuses on 'the activity of representing' (*ibid.*, p. 743). Analysing epistemic representation in these terms amounts to analysing schemes like ' S uses X to represent W for purposes P ' (*ibid.*, p. 743), or in more detail: 'Agents (1) intend; (2) to use model, M ; (3) to represent a part of the world W ; (4) for purposes, P . So agents specify which similarities are intended and for what purpose' (2010, p. 274). This conception of representation had already been proposed half century earlier by Apostel when he urged the following analysis of model-representation: 'Let then $R(S, P, M, T)$ indicate the main variables of the modelling relationship. The subject S takes, in view of the purpose P , the entity M as a model for the prototype T ' (Apostel 1961, p. 4). Including the intentions of model agents in the definition of scientific representation is now widely accepted, as we discuss in more detail in Section 5 (although Rusanen and Lappi disagree with this, and claim that 'the semantics of models as scientific representations should be based on the mind-independent model-world relation' (2012, p. 317)).

Giere's proposal, in our own terminology comes down to:

Similarity 5: A scientific model M represents a target system T iff there is an agent A who uses M to represent a target system T by proposing a theoretical hypothesis H specifying a similarity (in certain respects and to certain degrees) between M and T for purpose P .

This definition inherits from *Similarity 4* the resolutions of the problems of directionality, misrepresentation, and mistargeting; and for the sake of argument we assume that the problem with non-existent targets can be resolved in one way or other.

A crucial thing to note about *Similarity 5* is that, by invoking an active role for the purposes and actions of scientists in constituting epistemic representation, it marks a significant change in emphasis for similarity-based accounts. Suárez (2003, pp. 226-227), drawing on van Fraassen's (2002) and Putnam's (2002), defines 'naturalistic' accounts of representation as ones where 'whether or not representation obtains depends on facts about the world and does not in any way answer to the personal purposes, views or interests of enquirers'. By building the purposes of model users directly into an answer to the ER-problem, *Similarity 5* is explicitly not a naturalistic account (in contrast, for example, to *Similarity 1*). As noted in Section 2, we do not demand a naturalistic account of model-representation (and as we will see later, many of the more developed answers to the ER-problem are also not naturalistic accounts).

Does this suggest that *Similarity 5* is successful similarity-based solution to the ER-Problem? Unfortunately not. A closer look at *Similarity 5* reveals that the role of similarity has shifted. As far as offering a solution to the ER-Problem is concerned, all the heavy lifting in *Similarity 5* is done by the appeal to agents and similarity has in fact become an idle wheel. Giere implicitly admits this when he writes:

'How do scientists use models to represent aspects of the world? What is it about models that makes it possible to use them in this way? One way, perhaps the most important way, *but probably not the only way*, is by exploiting similarities between a model and that aspect of the world it is being used to represent. Note that I am not saying that the model itself represents an aspect of the world because it is similar to that aspect. There is no such representational relationship. [footnote omitted] Anything is similar to anything else in countless respects, but not anything represents anything else. *It is not the model that is doing the representing; it is the*

scientist using the model who is doing the representing.' (2004, p. 747, emphasis added)

But if similarity is not the only way in which a model can be used as a representation, and if it is the use by a scientist that turns a model into a representation (rather than any mind-independent relationship the model bears to the target), then similarity has become otiose in a reply to the ER-problem. A scientist could invoke any relation between M and T and M would still represent T . Being similar in the relevant respects to the relevant degrees now plays the role either of a representational style, or of a normative criterion for accurate representation, rather than of a grounding of representation. We assess in the next section whether similarity offers a cogent reply to the issues of style and accuracy.

A further problem is that there seems to be a hidden circularity in the analysis. As Toon (2012, pp. 251-252) points out, having a scientist form a theoretical hypothesis about the similarity relation between two objects A and B and exploit this similarity for a certain purpose P is not sufficient for representation. A and B could be two cars in a showroom and an engineer inspects car A and then use her knowledge about similarities to make assertions about B (for instance if both cars are of the same brand she can infer something about B 's quality of manufacturing). This, Toon submits, is not a case of representation: neither car is representational. Yet, if we delete the expression 'to represent' on the right hand side of the biconditional in *Similarity 5*, the resulting condition provides an accurate description of what happens in the showroom. So the only difference between the non-representational activity of comparing cars and representing B by A is that in one case A is 'used to represent' and in the other it's only 'used'. So representation is explained in terms of 'to represent', which is circular. So *Similarity 5* does not provide non-trivial conditions for something to be used *as a representation*.

One way around the problem would be to replace 'to represent' by 'to denote'. This, however, would bring the account close to *Similarity 3*, and it would suffer from the same problems.

Mäki (2011) suggested an extension of *Similarity 5*, which he explicitly brands as 'a (more or less explicit) version' of Giere's. Mäki adds two conditions to Giere's: the

agent uses the model to address an audience E and adds a commentary C (*ibid.*, p. 57). The role of the commentary is to specify the nature of the similarity. This is needed because ‘representation does not require that all parts of the model resemble the target in all or just any arbitrary respects, or that the issue of resemblance legitimately arises in regard to all parts. The relevant model parts and the relevant respects and degrees of resemblance must be delimited.’ (*ibid.*, p. 57). What these relevant respects and degrees of resemblance are depends on the purposes of the scientific representation in question. These are not determined ‘in the model’ as it were, but are pragmatic elements. From this it transpires that in effect C plays the same role as that played by theoretical hypotheses in Giere’s account. Certain aspects of M are chosen as those relevant to the representational relationship between M and T .

The addition of an audience, however, is problematic. While models are often shared publicly, this does not seem to be a necessary condition for the representational use of a model. There is nothing that precludes a lone scientist from coining a model M and using it representationally. That some models are easier to grasp, and therefore serve as more effective tools to drive home a point in certain public settings, is an undisputable fact, but one that has no bearing on a model’s status as a representation. The pragmatics of communication and the semantics of modelling are separate issues.

The conclusion we draw from this discussion is that similarity does not offer a viable answer to the ER-Problem.

4.2 Accuracy and Style

Accounting for the possibility of misrepresentation resulted in a shift of the division of labour for the more developed similarity based accounts. Rather than being the relation that grounds representation, similarity should be considered as setting a standard of accuracy or as providing an answer to the question of style (or both). The former is motivated by the observation that a proposed similarity between M and T could be wrong, and hence if the model user’s proposal does in fact hold (and M and T are in fact similar in the specified way) then M is an accurate representation of T . The latter transpires from the simple observation that a judgment of accuracy in fact

presupposes a choice of respects in which M and T are claimed to be similar. Simply proposing that they are similar in some unspecified respect is vacuous. But delineating relevant properties could potentially provide an answer to the problem of style. For example, if M and T are proposed to be similar with respect to their causal structure, then we might have a style of causal modelling; if M and T are proposed to be similar with respect to structural properties, then we might have a style of structural modelling; and so on and so forth. So the idea is that if M representing T involves the claim that M and T are similar in a certain respect, the respect chosen specifies the style of the representation; and if M and T are in fact similar in that respect (and to the specified degree), then M accurately represents T within that style.

In this section we investigate both options. But before delving into the details, let us briefly step back and reflect on possible constraints on viable answers. Taking his cue from Lopes' (2004) discussion of pictures, Downes (2009, pp. 421-422) proposes two constraints on allowable notions of similarity. The first, which he calls the *independence challenge*, requires that a user must be able to specify the relevant representation-grounding similarity *before* engaging a comparison between M and T . Similarities that are recognisable only with hindsight are an unsound foundation of a representation. We agree with this requirement, which in fact is also a consequence of the Surrogate Reasoning Condition: a model can generate novel hypotheses only if (at least some of the) similarity claims are not known only *ex post facto*.

Downes' second constraint, the *diversity constraint*, is the requirement that the relevant notion of similarity has to be identical in all kinds of representation and across all representational styles. So all models must bear the same similarity relations to their targets. Whatever its merits in the case of pictorial representation, this observation does not hold water in the case of scientific representation. Both Giere and Teller have insisted – rightly, in our view – that there need not be a substantive sense of similarity uniting all representations (see also Callender and Cohen(2006, p. 77) for a discussion). A proponent of the similarity view is free to propose different kinds of similarity for different representations and is under no obligation to also show that they are special cases of some overarching conception of similarity.

We now turn to the issue of style. A first step in the direction of an understanding of styles is the explicit analysis of the notion of similarity. Unfortunately the philosophical literature contains surprisingly little explicit discussion about what it means for something to be similar to something else. In many cases similarity is taken to be primitive, possible worlds semantics being a prime example. The problem is then compounded by the fact that the focus is on comparative overall similarity instead rather than on similarity in respect and degrees; for a critical discussion see (Morreau 2010). Where the issue is discussed explicitly, the standard way of cashing out what it means for an object to be similar to another object is to require that they co-instantiate properties. This is the idea that Quine (1969, pp. 117-118) and Goodman (1972, p. 443) had in mind in their influential critiques of the notion. They note that if all that is required for two things to be similar is that they co-instantiate *some* property, then everything is similar to everything else, since any pair of objects have at least one property in common.

The issue of similarity seems to have attracted more attention in psychology. In fact, the psychological literature provides formal accounts to capture it directly in more fully worked out accounts. The two most prominent suggestions are the *geometric* and *contrast* accounts (see (Decock and Douven 2011) for an up to date discussion). The former, associated with Shepard (1980), assigns objects a place in a multidimensional space based on values assigned to their properties. This space is then equipped with a metric and the degree of similarity between two objects is a function of the distance between the points representing the two objects in that space.

This account is based on the strong assumptions that values can be assigned to all features relevant to similarity judgments, which is deemed unrealistic. This problem is supposed to be overcome in Tversky's *contrast account* (1977). This account defines a gradated notion of similarity based on a weighted comparison of properties. Weisberg (2012; 2013, Ch. 8) has recently introduced this account into the philosophy of science where it serves as the starting point for his so-called *weighted feature matching account of model world-relations*. This account is our primary interest here.

The account introduces a set Δ of relevant properties. Let then $\Delta_M \subseteq \Delta$ be the set of properties from Δ that are instantiated by the model M ; likewise Δ_T is the set of properties from Δ instantiated by the target system. Furthermore let f be a ranking function assigning a real number to every subset of Δ . The simplest version of a ranking function is one that assigns to each set the number of properties in the set, but rankings can be more complex, for instance by giving important properties more weight. The level of similarity between M and T is then given by the following equation (Weisberg 2012, p. 788) (the notation is slightly amended):

$$S(M, T) = \theta f(\Delta_M \cap \Delta_T) - \alpha f(\Delta_M - \Delta_T) - \beta f(\Delta_T - \Delta_M),$$

where α, β and θ are weights, which can in principle take any value. This equation provides a ‘similarity score that can be used in comparative judgments of similarity’ (*ibid.*, p. 788). The score is determined by weighing the properties the model and target have in common against those they do not. (Thus we note that this account could be seen as a quantitative version of Hesse’s (1963) theory of analogy in which properties that M and T share are the *positive analogy* and ones they don’t share are the *negative analogy*). In the above formulation the similarity score S can in principle vary between any two values (depending on the choice of the ranking function and the value of the weights). One can then use standard mathematical techniques to renormalize S so that it takes values in the unit interval $[0,1]$ (these technical moves need not occupy us here and we refer the reader to Weisberg 2013, Chapter 8, for details).

The obvious question at this point is how the various blanks in the account can be filled. First in line is the specification of a property set Δ . Weisberg is explicit that there are no general rules to rely on and that ‘the elements of Δ come from a combination of context, conceptualization of the target, and theoretical goals of the scientist’ (Weisberg 2013, p. 149). Likewise, the ranking function as well as the values of weighting parameters depend on the goals of the investigation, the context, and the theoretical framework in which the scientists operate. Weisberg further divides the elements of Δ into *attributes* and *mechanisms*. The former are the ‘the properties and patterns of a system’ while the latter are the ‘underlying mechanism[s]

that generates these properties' (Weisberg 2013, p. 145). This distinction is helpful in the application to concrete cases, but for the purpose of our conceptual discussion it can be set aside.

Irrespective of these choices, the similarity score S has a number of interesting features. First, it is asymmetrical for $\alpha \neq \beta$, which makes room for the possibility of M being similar to T to a different degree than T is similar to M . So S provides the asymmetrical notion of similarity mentioned in subsection 4.1. Second, S has a property called *maximality*: everything is maximally similar to itself and every other non-identical object is equally or less similar. Formally: $S(A, A) \geq S(A, B)$ for all objects A and B as long as $A \neq B$ (*ibid.*, p. 154).

What does this account contribute to a response to the question of style? The answer, we think, is that it has heuristic value but does not provide substantive account. In fact, stylistic questions stand outside the proposed framework. The framework can be useful in bringing questions into focus, but eventually the substantive stylistic questions concern inclusion criteria for Δ (what properties do we focus on?), the weight given by f to properties (what is the relative importance of properties?) and the value of the parameters (how significant are disagreements between the properties of M and T ?). These questions have to be answered outside the account. The account is a framework in which questions can be asked but which does not itself provide answers, and hence no classification of representational styles emerges from it.

Some will say that this is old news. Goodman denounced similarity as 'a pretender, an impostor, a quack' (Goodman 1972, p. 437) not least because he thought that it merely put a label to something unknown without analysing it. And even some proponents of the similarity view have insisted that no general characterisation of similarity was possible. Thus Teller submits that '[t]here can be no general account of similarity, but there is also no need for a general account because the details of any case will provide the information which will establish just what should count as relevant similarity in that case.' (2001, p. 402) This amounts to nothing less than the admission that no analysis of similarity (or even different kinds of similarity) is possible and that we have to deal with each case in its own right.

Assume now, for the sake of argument, that the stylistic issues have been resolved and full specifications of relevant properties and their relative weights are available. It would then seem plausible to say that $S(M, T)$ provides a degree of accuracy. This reading is supported by the fact that Weisberg paraphrases the role of $S(M, T)$ as providing ‘standards of fidelity’ (Weisberg 2013, p. 147). Indeed, in response to Parker’s (2015), Weisberg claims that his weighted feature matching account is supposed to answer the ER-problem *and* provide standards of accuracy.

As we have seen above, $S(M, T)$ is maximal if M is a perfect replica of T (with respect to the properties in Δ), and the fewer properties M and T share, the less accurate the representation becomes. This lack of accuracy is then reflected in a lower similarity score. This is plausible and Weisberg’s account is indeed a step forward in the direction of quantifying accuracy.

Weisberg’s account is an elaborate version of the co-instantiation account of similarity. It improves significantly on simple versions, but it cannot overcome that account’s basic limitations. Niiniluoto distinguishes between two different kinds of similarities (Niiniluoto 1988, pp. 272-274): partial identity and likeness (which also feature in Hesse’s discussion of analogies, see, for instance (1963, pp. 66-67)). Assume M instantiates the relevant properties P_1, \dots, P_n and T instantiates the relevant properties Q_1, \dots, Q_n . If these properties are identical, i.e. if $P_i = Q_i$ for all $i = 1, \dots, n$, then M and T are similar in the sense of being *partially identical*. Partial identity contrasts with what Niiniluoto calls *likeness*. M and T are similar in the sense of likeness if the properties are not identical but similar themselves: P_i is similar to Q_i for all $i = 1, \dots, n$. So in likeness the similarity is located at the level of the properties themselves. For example, a red post box and a red London bus are similar with respect to their colour, even if they do not instantiate the exact same shade of red. As Parker (2015’, p. 273) notes, Weisberg’s account (like all co-instantiation accounts) deals well with partial identity, but has no systematic place for likeness. To deal with likeness Weisberg would in effect have to reduce likeness to partial identity by introducing ‘imprecise’ properties which encompass the P_i and the Q_i . Parker (*ibid.*) suggest that this can be done by introducing intervals in the feature set, for instance of

the form ‘the value of feature X lies in the interval $[x - \varepsilon, x + \varepsilon]$ ’ where ε is parameter specifying the precision of overlap. To illustrate she uses Weisberg’s example of the San Francisco Bay model and claims that in order to account for the similarity between the model and the actual Bay with respect to their Froude number Weisberg has to claim something like: ‘The Bay model and the real Bay share the property of having a Froude number that is within 0.1 of the real Bay’s number. It is more natural to say that the Bay model and the real Bay have *similar* Froude numbers—similar in the sense that their values differ by at most 0.1.’ (Parker 2015, p. 273)

In his response Weisberg accepts this and argues that he is trying to provide a reductive account of similarity that bottoms out in properties shared and those not shared (2015, p. 302). But such interval-valued properties have to be part of Δ in order for the formal account to capture them. This means that another important decision regarding whether or not M and T are similar occurs outside of the formal account itself. The inclusion criteria on what goes into Δ now not only has to delineate relevant properties, but, at least for the quantitative ones, also has to provide an interval defining when they qualify as similar. Furthermore, it remains unclear how to account for M and T to be alike with respect to their qualitative properties. The similarity between genuinely qualitative properties cannot be accounted for in terms of numerical intervals. This is a particularly pressing problem for Weisberg, because he takes the ability to compare models and their targets with respect to their qualitative properties as a central desideratum for any account of similarity between the two (Weisberg 2013, p. 136).

4.3 Problems of Ontology

Another problem facing similarity based approaches concerns their treatment of the ontology of models. If models are supposed to be similar to their targets in the ways specified by theoretical hypotheses or commentaries, then they must be the *kind* of things that can be so similar.

Some models are homely physical objects. The Army Corps of Engineers' model of the San Francisco Bay is a water basin and equipped with pumps to simulate the action of tidal flows (Weisberg 2013); ball and stick models of molecules are made of metal or wood (Toon 2011); the Phillips-Newlyn model of an economy is system of pipes and reservoirs (Morgan and Knuuttila 2012); and model organisms in biology are animals like worms and mice (Ankeny and Leonelli 2011). For models of this kind similarity is straightforward (at least in principle) because they are of the same ontological kind as their respective targets: they are material objects.

But many interesting scientific models are not like this. Two perfect spheres with a homogeneous mass distribution which interact only with each other (the Newtonian model of the sun-earth system) or a single-species population isolated from its environment and reproducing at fixed rate at equidistant time steps (the logistic growth model of a population) are what Hacking aptly describes as 'something you hold in your head rather than your hands' (1983, p. 216). Following Thomson-Jones (2012) we call such models *non-concrete models*. The question then is what kind of objects non-concrete models are. Giere submits that they are abstract objects: 'Models in advanced sciences such as physics and biology should be abstract objects constructed in conformity with appropriate general principles and specific conditions.' ((Giere 2004, p. 747) *cf.* (Giere 1988, p. 81; 2010, p. 270)).

The appeal to abstract entities brings a number of difficulties with it. The first is that the class of abstract objects is rather large. Numbers and other objects of pure mathematics, classes, propositions, concepts, the letter 'A', and Dante's *Inferno* are abstract objects (Rosen 2014), and Hale (1988, pp. 86-87) lists no less than 12 different possible characterisations of abstract objects. At the very least this list shows that there is great variety in abstract objects and classifying models as abstract object adds little specificity to an account of what models are. Giere could counter that he limits attention to those abstract objects that possess 'all and only the characteristics specified in the principles' (Giere 2004, p. 745), where principles are general rules like Newton's laws of motion. He further specifies that he takes 'abstract entities to be human constructions' and that 'abstract models are definitely not to be identified with linguistic entities such as words or equations' (*ibid.*, p. 747). While this narrows down

the choices somehow, it still leaves many options and ultimately the ontological status of models in a similarity account remains unclear.

Giere fails to expand on this ontological issue for a reason: he dismisses the problem as one that philosophers of science can set aside without loss. He voices scepticism about the view that philosophers of science ‘need a deeper understanding of imaginative processes and of the objects produced by these process’ (Giere 2009, p.250) or that ‘we need say much more [...] to get on with the job of investigating the functions of models in science’ (*ibid.*).

We remain unconvinced about this scepticism, not least because there is an obvious yet fundamental issue with abstract objects. No matter how the above issues are resolved (and irrespective of whether they are resolved at all), at the minimum it is clear that models are ‘abstract’ in the sense that they have no spatiotemporal location. Teller (2001, p. 399) and Thomson-Jones (2010) supply arguments suggesting that this alone causes serious problems for the similarity account. The similarity account demands that models can instantiate properties and relations, since this is a necessary condition on them being similar to their targets. In particular, it requires that models can instantiate the properties and relations mentioned in theoretical hypotheses or commentaries. But such properties and relations are typically *physical*. And if models have no spatial-temporal location, then they do not instantiate any such properties or relations. Thomson-Jones’ example of the idealized pendulum model makes this clear. If the idealized pendulum is abstract then it is difficult to see how to make sense of the idea that it has a length, or a mass, or an oscillation period of any particular time.

An alternative suggestion due to Teller (2001) is that we should instead say that whilst ‘concrete objects HAVE properties [...] properties are PARTS of models’ (*ibid.* p. 399, original capitalisation). It is not entirely clear what Teller means by this, but our guess is that he would regard models as bundles of properties. Target systems, as concrete objects, are the sorts of things that can instantiate properties delineated by theoretical hypotheses. Models, since they are abstract, cannot. But rather than being objects instantiating properties, a model can be seen as a bundle of properties. A collection of properties is an abstract entity that is the sort of thing that can contain

the properties specified by theoretical hypotheses as parts. The similarity relation between models and their targets shifts from the co-instantiation of properties, to the idea that targets instantiate (relevant) properties that are parts of the model. With respect to what it means for a model to be a bundle of properties Teller claims that the '[d]etails will vary with ones account of instantiation, of properties and other abstract objects, and of the way properties enter into models' (*ibid*).

But as Thompson-Jones (2010, pp. 294-295) notes, it is not obvious that this suggestion is an improvement on Giere's abstract objects. A bundle view incurs certain metaphysical commitments, chiefly the existence of properties and their abstractness, and a bundle view of objects, concrete or abstract, faces a number of serious problems (Armstrong 1989). One might speculate that addressing these issues would push Teller either towards the kind of more robust account of abstract objects that he endeavoured to avoid, or towards a fictionalist understanding of models.

The latter option has been discussed by Giere, who points out that a natural response to Teller's and Thomson-Jones' problem is to regard models as akin to *imaginary* or *fictional* systems of the sort presented in novels and films. It seems true to say that Sherlock is a smoker, despite the fact that Sherlock an imaginary detective, and smoking is a physical property. At times, Giere seems sympathetic to this view. He says 'it is widely assumed that a work of fiction is a creation of human imagination ... the same is true of scientific models. So, ontologically, scientific models and works of fiction are on a par. They are both imaginary constructs.' (2009, p. 249), and he observes that 'novels are commonly regarded as works of imagination. That, ontologically, is how we should think of abstract scientific models. They are creations of scientists' imaginations. They have no ontological status beyond that.' (2010, p. 278) However, these seem to be occasional slips and he recently positioned himself as an outspoken opponent of any approach to models that likens them to literary fiction. We discuss these approaches as well as Giere's criticisms of them in Section 7.

In sum, the similarity view is yet to be equipped with a satisfactory account of the ontology of models.

5. The Structuralist Conception

The structuralist conception of model-representation originated in the so-called semantic view of theories that came to prominence in the second half of the 20th century (Suppes (2002), van Fraassen (1980), and Da Costa and French (2003) provide classical statements of the view; Byerly (1969), Chakravartty (2001), Klein (2013) and Portides (2005; 2010) provide critical discussions). The semantic view was originally proposed as an account of theory structure rather than model-representation. The driving idea behind the position is that scientific theories are best thought of as collections of models. This invites the questions: what are these models, and how do they represent their target systems? Defenders of the semantic view of theories take models to be structures, which represent their target systems in virtue of there being some kind of *mapping* (isomorphism, partial isomorphism, homomorphism, ...) between the two. (It is worth noting that Giere, whose account of scientific representation we discussed in the previous section, is also associated with the semantic view, despite not subscribing to either of these positions.)

This conception has two *prima facie* advantages. The first advantage is that it offers a straightforward answer to the ER-Problem, and one that accounts for surrogate reasoning: the mappings between the model and the target allow scientists to convert truths found in the model into claims about the target system. The second advantage concerns the applicability of mathematics. There is a time-honoured position in the philosophy of mathematics which sees mathematics as the study of structures; see, for instance, Resnik's (1997) and Shapiro's (2000). It is a natural move for the scientific structuralist to adopt this point of view, which, without further ado, provides a neat explanation of how mathematics is used in scientific modelling.

5.1 Structures and the Problem of Ontology

Almost anything from a concert hall to a kinship system can be referred to as a 'structure'. So the first task for a structuralist account of representation is to articulate what notion of structure it employs. A number of different notions of structure have been discussed in the literature (for a review see Thomson-Jones' (2011)), but by far

the most common and widely used is the notion of structure one finds in set theory and mathematical logic. A structure S in that sense (sometimes ‘mathematical structure’ or ‘set-theoretic structure’) is a composite entity consisting of the following: a non-empty set U of objects called the domain (or universe) of the structure and a non-empty indexed set R of relations on U . With the exception of the caveat below regarding interpretation functions, this definition of structure is widely used in mathematics and logic; see for instance Machover’s (1996, p. 149), Hodges’ (1997, p. 2), and Rickart’s (1995, p. 17)). It is convenient to write these as $S = \langle U, R \rangle$, where ‘ \langle, \rangle ’ denotes an ordered tuple. Sometimes operations are also included in the definition of a structure. While convenient in some applications, operations are redundant because operations reduce to relations (see Boolos and Jeffrey’s (Boolos and Jeffrey 1989, pp. 98-99)).

It is important to be clear on what we mean by ‘object’ and ‘relation’ in this context. As Russell (1919/1993, p. 60) points out, in defining the domain of a structure it is irrelevant what the objects are. All that matters from a structuralist point of view is that there are so and so many of them. Whether the object is a desk or a planet is irrelevant. All we need are dummies or placeholders whose *only* property is ‘objecthood’. Similarly, when defining relations one disregards completely what the relation is ‘in itself’. Whether we talk about ‘being the mother of’ or ‘standing to the left of’ is of no concern in the context of a structure; all that matters is between which objects it holds. For this reason, a relation is specified purely extensionally: as a class of ordered n -tuples. The relation literally is nothing over and above this class. So a structure consists of dummy-objects between which purely extensionally defined relations hold.

Let us illustrate this with an example. Consider the structure with the domain $U = \{a, b, c\}$ and the following two relations: $r_1 = \{a\}$ and $r_2 = \{\langle a, b \rangle, \langle b, c \rangle, \langle a, c \rangle\}$. Hence R consists of r_1 and r_2 , and the structure itself is $S = \langle U, R \rangle$. This is a structure with a three-object domain endowed with a monadic property and a transitive relation. Whether the objects are books or iron rods is of no relevance to the structure; they could be literally anything one can think of. Likewise r_1 could be

literally any monadic property (being green, being waterproof, etc) and r_2 could be any (irreflexive) transitive relation (larger than, hotter than, more expensive than, etc.).

It is worth pointing out that this use of ‘structure’ differs from the use one sometimes finds in logic, where linguistic elements are considered part of the model as well. Specifically, over and above $S = \langle U, R \rangle$, a structure is also taken to include a language (sometimes called a ‘signature’) L , and an interpretation function (see, for example, (Enderton 1972/2001, pp. 80-81; Hodges 1997, Ch.1)). But in the context of the accounts discussed in this section, a structure is the ordered pair $S = \langle U, R \rangle$ as introduced above and so we disregard this alternative use of ‘structure’.

The first basic posit of the structuralist theory of representation is that models are structures in the above sense (the second is that models represent their targets by being suitably morphic to them; we discuss morphisms in the next subsection). Suppes has articulated this stance clearly when he declared that ‘the meaning of the concept of model is the same in mathematics and the empirical sciences’ (1960, p. 12). Likewise, van Fraassen posits that a ‘scientific theory gives us a family of models to represent the phenomena’, that ‘[t]hese models are mathematical entities, so all they have is structure [...]’ (1997, pp. 528-529) and that therefore ‘[s]cience is [...] interpreted as saying that the entities stand in relations which are transitive, reflexive, etc. but as giving no further clue as to what those relations are’ (1997, p. 516). Redhead submits that ‘it is this abstract structure associated with physical reality that science aims, and to some extent succeeds, to uncover [...]’ (2001, p. 75). Finally, French and Ladyman affirm that ‘the specific material of the models is irrelevant; rather it is the structural representation [...] which is important’ (1999, p. 109). Further explicit statements of this view are offered by: Da Costa and French (1990, p. 249), Suppes, (1962/1969, p. 24; 1970, Ch.2) and van Fraassen (1980, pp. 43,64; 1991, p. 483; 1995, p. 6; 1997, p. 516, 522).

These structuralist accounts have typically been proposed in the framework of the so-called semantic view of theories. There are differences between them, and formulations vary from author to author. However, as Da Costa and French (2000)

point out, all these accounts share a commitment to analysing models as structures. So we are presented with a clear answer to the Problem of Ontology: models are structures. The remaining issue is what structures themselves are. Are they Platonic entities, equivalence classes, modal constructs, or yet something else? This is a hotly debated issue in the philosophy of logic and mathematics; for different positions see for instance Dummett's (1991, 295ff.), Hellman's (1989; 1996), Redhead's (2001), Resnik's (1997), and Shapiro's (2000). But philosophers of science need not resolve this issue and can pass off the burden of explanation to philosophers of mathematics. This is what usually happens, and hence we don't pursue this matter further.

An extension of the standard conception of structure is the so-called partial structures approach (see, for instance, Da Costa and French's (2003) and Bueno, French, and Ladyman's (2002)). Above we defined relations by specifying between which tuples it holds. This naturally allows a sorting of all tuples into two classes: ones that belong to the relation and ones that don't. The leading idea of partial structures is to introduce a third option: for some tuples it is indeterminate whether or not they belong to the relation. Such a relation is a *partial relation*. A structure with a set R containing partial relations is a *partial structure* (formal definitions can be found in references given above). Partial structures make room for a process of scientific investigation where one begins not knowing whether a tuple falls under the relation and then learns whether or not it does.

Proponents of that approach are more guarded as regards the ontology of models. Bueno and French emphasise that 'advocates of the semantic account need not be committed to the ontological claim that models *are* structures' (2011, p. 890 original emphasis). This claim is motivated by the idea that the task for philosophers of science is to represent scientific theories and models, rather than to reason about them directly. French (2010) makes it explicit that according to his account of the semantic view of theories, a scientific theory is *represented* as a class of models, but should not be identified with that class. Moreover, a class of models is just one way of representing a theory; we can also use an intrinsic characterisation and represent the same theory as a set of sentences in order to account for how they can be objects of our epistemic attitudes (French and Saatsi 2006).

He therefore adopts a quietist position with respect to what a theory or a model *is*, declining to answer the question (2010; cf. French and Vickers 2011). There are thus two important notions of representation at play: representation of targets by models, which is the job of scientists, and representation of theories and models by structures, which is the job of philosophers of science. The question for this approach then becomes whether or not the structuralist representation of models and epistemic representation – as partial structures and morphisms that hold between them – is an accurate or useful one. And the concerns raised below remain when translated into this context as well.

There is an additional question regarding the correct formal framework for thinking about models in the structuralist position. Landry (2007) argues that in certain contexts group, rather than set, theory should be used when talking about structures and morphisms between them, and Halvorson (2012; forthcoming) argues that theories should be identified with categories rather than classes or sets. Although these discussions highlight important questions regarding the nature of scientific theories, the question of how individual models represent remains unchanged. Halvorson still takes individual models to be set-theoretic structures. And Landry's paper is not an attempt to reframe the representational relationship between models and their targets (see (Brading and Landry 2006) for her scepticism regarding how structuralism deals with this question). Thus, for reasons of simplicity we will focus on the structuralist view that identifies models with set-theoretic structures throughout the rest of this section.

5.2 Structuralism and the ER-Problem

The most basic structuralist conception of scientific representation asserts that scientific models, understood as structures, represent their target systems in virtue of being isomorphic to them. Two structures $S_a = \langle U_a, R_a \rangle$ and $S_b = \langle U_b, R_b \rangle$ are isomorphic iff there is a mapping $f : U_a \rightarrow U_b$ such that (i) f is one-to-one (bijective) and (ii) f preserves the system of relations in the following sense: the members a_1, \dots, a_n of U_a satisfy the relation r_a of R_a iff the corresponding members $b_1 = f(a_1), \dots, b_n = f(a_n)$ of U_b satisfy the relation r_b of R_b , where r_b is the relation

corresponding to r_a (for difficulties in how to cash out this notion of correspondence without reference to an interpretation function see Halvorson's (2012) and Glymour's (2013)).

Assume now that the target system T exhibits the structure $S_T = \langle U_T, R_T \rangle$ and the model is the structure $S_M = \langle U_M, R_M \rangle$. Then the model represents the target iff it is isomorphic to the target:

Structuralism 1: A scientific model M represents its target T iff S_M is isomorphic to S_T .

This view is articulated explicitly by Ubbink, who posits that 'a model represents an object or matter of fact in virtue of this structure; so an object is a model [...] of matters of fact if, and only if, their structures are isomorphic' (1960, p. 302). Views similar to Ubbink's seem operable in many versions of the semantic view. In fairness to proponents of the semantic view it ought to be pointed out, though, that for a long time representation was not the focus of attention in the view and the attribution of (something like) *Structuralism 1* to the semantic view is an extrapolation. Representation became a much-debated topic in the first decade of the 21st century, and many proponents of the semantic view then either moved away from *Structuralism 1*, or pointed out that they never held such a view. We turn to more advanced positions shortly, but to understand what motivates such positions it is helpful to understand why *Structuralism 1* fails.

An immediate question concerns the target end structure S_T . At least prima facie target systems aren't structures; they are physical objects like planets, molecules, bacteria, tectonic plates, and populations of organisms. An early recognition that the relation between targets and structures is not straightforward can be found in Byerly, who emphasizes that structures are abstracted from objects (1969, p.135-138). The relation between structures and physical targets is indeed a serious question and we will return to it in subsection 5.4. In this subsection we grant the structuralist the assumption that target systems are (or at least have) structures.

The first and most obvious problem is the same as with the similarity view: isomorphism is symmetrical, reflexive, and transitive, but epistemic representation isn't. This problem could be addressed by replacing isomorphism with an alternative mapping. Bartels (2006), Lloyd (1984), and Mundy (1986) suggest homomorphism; van Fraassen (1980; 1997; 2008) and Redhead isomorphic embeddings (2001); advocates of the partial structures approach prefer partial isomorphisms (1997; 1999; 1990; 2003; 2000; 1999); and Swoyer (1991) introduces what he calls Δ/Ψ -morphisms. We refer to these collectively as 'morphisms'.

This solves some, but not all problems. While many of these mappings are asymmetrical, they are all still reflexive, and at least some of them are also transitive. But even if these formal issues could be resolved in one way or another, a view based on structural mappings would still face other serious problems. For ease of presentation we discuss these problems in the context of the isomorphism view; *mutatis mutandis* other formal mappings suffer from the same difficulties. (For detailed discussions of homomorphism and partial isomorphism see Suárez's (2003, pp. 239-241) and Pero and Suárez's (2016); Mundy (1986) discusses general constraints one may want to impose on morphisms.)

Like similarity, isomorphism is too inclusive: not all things that are isomorphic represent each other. In the case of similarity this case was brought home by Putnam's thought experiment with the ant crawling on the beach; in the case of isomorphism a look at the history of science will do the job. Many mathematical structures have been discovered and discussed long before they have been used in science. Non-Euclidean geometries were studied by mathematicians long before Einstein used them in the context of spacetime theories, and Hilbert spaces were studied by mathematicians prior to their use in quantum theory. If representation was nothing over and above isomorphism, then we would have to conclude that Riemann discovered general relativity or that Hilbert invented quantum mechanics. This is obviously wrong. Isomorphism on its own does not establish representation (Frigg 2002, p. 10).

Isomorphism is more restrictive than similarity: not everything is isomorphic to everything else. But isomorphism is still too abundant to correctly identify the extension of a representation (i.e. the class of systems it represents), which gives rise to a version of the mistargeting problem. The root of the difficulties is that the same structures can be instantiated in different target systems. The $1/r^2$ law of Newtonian gravity is also the ‘mathematical skeleton’ of Coulomb’s law of electrostatic attraction and the weakening of sound or light as a function of the distance to the source. The mathematical structure of the pendulum is also the structure of an electric circuit with condenser and solenoid (a detailed discussion of this case is provided by Kroes (1989)). Linear equations are ubiquitous in physics, economics and psychology. Certain geometrical structures are instantiated by many different systems; just think about how many spherical things we find in the world. This shows that the same structure can be exhibited by more than one target system. Borrowing a term from the philosophy of mind, one can say that structures are *multiply realisable*. If representation is explicated solely in terms of isomorphism, then we have to conclude that, say, a model of a pendulum also represents an electric circuit. But this seems wrong. Hence isomorphism is too inclusive to correctly identify a representation’s extension.

One might try to dismiss this point as an artefact of a misidentification of the target. Van Fraassen (1980, p. 66), mentions a similar problem under the heading of ‘unintended realisations’ and then expresses confidence that it will ‘disappear when we look at larger observable parts of the world’. Even if there are multiply realisable structures to begin with, they vanish as science progresses and considers more complex systems because these systems are unlikely to have the same structure. Once we focus on a sufficiently large part of the world, no two phenomena will have the same structure. There is a problem with this counter, however. To appeal to *future* science to explain how models work *today* seems unconvincing. It is a matter of fact that we *currently* have models that represent electric circuits and sound waves, and we do not have to await future science providing us with more detailed accounts of a phenomenon to make our models represent what they actually already do represent.

As we have seen in the last section, a misrepresentation is one that portrays its target as having features it doesn’t have. In the case of an isomorphism account of

representation this presumably means that the model portrays the target as having structural properties that it doesn't have. However, isomorphism demands identity of structure: the structural properties of the model and the target must correspond to one another exactly. A misrepresentation won't be isomorphic to the target. By the lights of *Structuralism I* it is therefore is not a representation at all. Like simple similarity accounts, *Structuralism I* conflates misrepresentation with non-representation.

Muller (2011, p. 112) suggests that this problem can be overcome in a two-stage process: one first identifies a sub-model of the model, which in fact is isomorphic to at least a part of the target. This 'reduced' isomorphism establishes representation. One then constructs a 'tailor-made morphism on a case by case basis' (*ibid.*, p. 112) to account for accurate representation. Muller is explicit that this suggestion presupposes that there is 'at least one resemblance' (*ibid.*, p. 112) between model and target because 'otherwise one would never be called a representation of the other' (*ibid.*, p. 112). While this may work in some cases, it is not a general solution. It is not clear whether all misrepresentations have isomorphic sub-models. Models that are gross distortions of their targets (such as the liquid drop model of the nucleus or the logistic model of a population) may well not have such sub-models. More generally, as Muller admits, his solution 'precludes total misrepresentation' (*ibid.*, p. 112). So in effect it just limits the view that representation coincides with correct representation to a sub-model. However, this is too restrictive a view of representation. Total misrepresentations may be useless, but they are representations nevertheless.

Another response refers to the partial structures approach and emphasises that partial structures are in fact constructed to accommodate a mismatch between model and target and are therefore not open to this objection (Bueno and French 2011, p. 888). It is true that the partial structures framework has a degree of flexibility that the standard view does not. However, we doubt that this flexibility stretches far enough. While the partial structure approach deals successfully with incomplete representations, it does not seem to deal well with distortive representations (we come back to this point in the next subsection). So the partial structures approach, while enjoying an advantage over the standard approach, is nevertheless not yet home and dry.

Like the similarity account, *Structuralism 1* has a problem with non-existent targets because no model can be isomorphic to something that doesn't exist. If there is no ether, a model can't be isomorphic to it. Hence models without target cannot represent what they seem to represent.

Most of these problems can be resolved by making moves similar to the ones that lead to *Similarity 5*: introduce agents and hypothetical reasoning into the account of representation. Going through the motions one finds:

Structuralism 2: A scientific model M represents a target system T iff there is an agent A who uses M to represent a target system T by proposing a theoretical hypothesis H specifying an isomorphism between S_M and S_T .

Something similar to this was suggested by Adams (1959, p. 259) who appeals to the idea that physical systems are the *intended* models of a theory in order to differentiate them from purely mathematical models of a theory. This suggestion is also in line with van Fraassen's recent pronouncements on representation. He offers the following as the 'Hauptstutz' of a theory of representation: '*There is no representation except in the sense that some things are used, made, or taken, to represent things as thus and so*' (2008, p. 23, original emphasis). Likewise, Bueno submits that 'representation is an *intentional* act relating two objects' (2010, p. 94, original emphasis), and Bueno and French point out that using one thing to represent another thing is not only a function of (partial) isomorphism but also depends on 'pragmatic' factors 'having to do with the use to which we put the relevant models' (2011, p. 885). This, of course, gives up on the idea of an account which reduces representation to intrinsic features of models and their targets. At least one extra element, the model user, also features in whatever relation is supposed to constitute the representational relationship between M and T . In a world with no agents, there would be no scientific representation.

This seems to be the right move. Like *Similarity 5*, *Structuralism 2* accounts for the directionality of representation and has no problem with misrepresentation. But, again as in the case of *Similarity 5*, this is Pyrrhic victory as the role of isomorphism has shifted. The crucial ingredient is the agent's intention and isomorphism has in fact

become either a representational style or normative criterion for accurate representation. Let us now assess how well isomorphism fares as a response to these problems.

5.3 Accuracy, Style and Demarcation

The problem of style is to identify representational styles and characterise them. Isomorphism offers an obvious response to this challenge: one can represent a system by coming up with a model that is structurally isomorphic to it. We call this the isomorphism-style. This style also offers a clear-cut condition of accuracy: the representation is accurate if the hypothesised isomorphism holds; it is inaccurate if it doesn't.

This is neat answer. The question is what status it has vis-à-vis the problem of style. Is the isomorphism-style merely one style among many other styles which are yet to be identified, or is it in some sense privileged? The former is uncontentious. However, the emphasis many structuralists place on isomorphism suggests that they do not regard isomorphism as merely one way among others to represent something. What they seem to have in mind is the stronger claim that a representation *must* be of that sort, or that the isomorphism-style is the only acceptable style.

This claim seems to conflict with scientific practice. Many representations are inaccurate in some way. As we have seen above, partial structures are well equipped to deal with incomplete representations. However, not all inaccuracies are due to something being left out. Some models distort, deform and twist properties of the target in ways that seem to undercut isomorphism. Some models in statistical mechanics have an infinite number of particles and the Newtonian model of the solar system represents the sun as perfect sphere where it in reality is fiery ball with no well-defined surface at all. It is at best unclear how isomorphism, partial or otherwise, can account for these kinds of idealisations. From an isomorphism perspective all one can say about such idealisations is that they are failed isomorphism representations (or isomorphism misrepresentations). This is rather uninformative. One might try to characterise these idealisations by looking at *how* they fail to be isomorphic to their targets, but we doubt that this is going very far. Understanding how distortive

idealisations work requires a positive characterisation of them, and we cannot see how such a characterisation could be given within the isomorphism framework. So one has to recognise styles of representation other than isomorphism.

This raises that question whether other mappings such as homomorphisms, or embeddings would fit the bill. They would, we think, make valuable additions to the list of styles, but they would not fill all gaps. Like isomorphism, these mappings are not designed to accommodate distortive idealisations, and hence a list of styles that includes them still remains incomplete.

Structuralism's stand on the demarcation problem is by and large an open question. Unlike similarity, which has been widely discussed across different domains, isomorphism is tied closely to the formal framework of set theory, and it has been discussed only sparingly outside the context of the mathematized sciences. An exception is French, who discusses isomorphism accounts in the context of pictorial representation (2003). He discusses in detail Budd's (1993) account of pictorial representation and points out that it is based on the notion of a structural isomorphism between the structure of the surface of the painting and the structure of the relevant visual field. Therefore representation is the perceived isomorphism of structure (French 2003, pp. 1475-1476) (this point is reaffirmed by Bueno and French (2011, pp. 864-865); see Downes's(2009, pp. 423-425) for a critical discussion). In a similar vein, Bueno claims that the partial structures approach offers a framework in which different representations – among them 'outputs of various instruments, micrographs, templates, diagrams, and a variety of other items' (2010, p. 94) – can be accommodated. This would suggest that an isomorphism account of representation at least has a claim to being a universal account covering representations across different domains.

This approach faces a number of questions. First, neither a visual field nor a painting is a structure, and the notion of there being an isomorphism in the set theoretic sense between the two at the very least needs unpacking. The theory is committed to the claim that paintings and visual fields have structures, but, as we will see in the next subsection, this claim faces serious issues. Second, Budd's theory is only one among many theories of pictorial representation, and most alternatives do not invoke

isomorphism. So there is question whether a universal claim can be built on Budd's theory. In fact, there is even a question about isomorphism's universality within scientific representation. Non-mathematised sciences work with models that aren't structures. Godfrey-Smith (2006), for instance, argues that models in many parts of biology are imagined concrete objects. There is a question whether isomorphism can explain how models of that kind represent.

This points to a larger issue. The structuralist view is a rational reconstruction of scientific modelling, and as such it has some distance from the actual practice. Some philosophers have worried that this distance is too large and that the view is too far removed from the actual practice of science to be able to capture what matters to the practice of modelling (this is the thrust of many contributions to (Morgan and Morrison 1999); see also (Cartwright 1999)). Although some models used by scientists may be best thought of as set theoretic structures, there are many where this seems to contradict how scientists actually talk about, and reason with, their models. Obvious examples include physical models like the San Francisco Bay model (Weisberg 2013), but also systems such as the idealized pendulum or imaginary populations of interbreeding animals. Such models have the strange property of being *concrete-if-real* and scientists talk about them as if they were real systems, despite the fact that they are obviously not. Thomson-Jones (2010) dubs this 'face value practice', and there is a question whether structuralism can account for that practice.

5.4 The structure of target systems

Target systems are physical objects: atoms, planets, populations of rabbits, economic agents, etc. Isomorphism is a relation that holds between two structures and claiming that a set theoretic structure is isomorphic to a piece of the physical world is *prima facie* a category mistake. By definition, all of the mappings suggested – isomorphism, partial isomorphism, homomorphism, or isomorphic embedding – only hold between two structures. If we are to make sense of the claim that the model is isomorphic to its target we have to assume that the target somehow exhibits a certain structure $S_T = \langle U_T, R_T \rangle$. But what does it mean for a target system – a part of the physical world – to possess a structure, and where in the target system is the structure located?

The two prominent suggestions in the literature are that data-models are the target end structures represented by models, and that structures are, in some sense, instantiated in target systems. The latter option comes in three versions. The first version is that a structure is ascribed to a system; the second version is that systems instantiate structural universals; and the third version claims that target systems simply are structures. We consider all suggestions in turn.

What are data models? Data are what we gather in experiments. When observing the motion of the moon, for instance, we choose a coordinate system and observe the position of the moon in this coordinate system at consecutive instants of time. We then write down these observations. The data thus gathered are called the *raw* data. The raw data then undergo a process of cleansing, rectification and regimentation: we throw away data points that are obviously faulty, take into consideration what the measurement errors are, take averages, and usually idealise the data, for instance by replacing discrete data points by a continuous function. Often, although not always, the result is a smooth curve through the data points that satisfies certain theoretical desiderata (Harris (2003) and van Fraassen (2008, pp. 166-68) elaborate on this process). These resulting data models can be treated as set theoretic structures. In many cases the data points are numeric and the data model is a smooth curve through these points. Such a curve is a relation over \mathbb{R}^n (for some n), or subsets thereof, and hence it is structure in the requisite sense.

Suppes (1962/1969) was the first to suggest that data models are the targets of scientific models: models don't represent parts of the world; they represent data structures. This approach has then been adopted by van Fraassen, when he declares that '[t]he whole point of having theoretical models is that they should fit the phenomena, that is, fit the models of data' (1981, p. 667). He has defended this position numerous times over the years (1980, p. 64; 1985, p. 271; 1989, p. 229; 1997, p. 524; 2002, p. 164) including in his most recent book on representation (2008, p. 246, 252). So models don't represent planets, atoms or populations; they represent data that are gathered when performing measurements on planets, atoms or populations.

This revisionary point of view has met with stiff resistance. Muller articulates the unease about this position as follows: ‘the best one could say is that a data structure D seems to act as *simulacrum* of the concrete actual being B [...] But this is not good enough. We don’t want simulacra. We want the real thing. Come on.’ (2011, p. 98). Muller’s point is that science aims (or at least has to aim) to represent real systems in the world and not data structures. Van Fraassen calls this the ‘loss of reality objection’ (2008, p. 258) and accepts that the structuralist must ensure that models represent target systems, rather than finishing the story at the level of data. In his (2008) he addresses this issue in detail and offers a solution. We discuss his solution below, but before doing so we want to articulate the objection in more detail. To this end we briefly revisit the discussion about phenomena and data which took place in the 1980s and 1990s.

Bogen and Woodward (1988), Woodward (1989), and more recently (and in a somewhat different guise) Teller (2001), introduced the distinction between phenomena and data and argue that models represent phenomena, not data. The difference is best introduced with an example: the discovery of weak neutral currents (Bogen and Woodward 1988, pp. 315-318). What the model at stake consists of is particles: neutrinos, nucleons, and the Z^0 particle, along with the reactions that take place between them. (The model we are talking about here is not the so-called standard model of elementary particles as a whole. Rather, what we have in mind is one specific model about the interaction of certain particles of the kind one would find in a theoretical paper on this experiment.) Nothing of that, however, shows in the relevant data. CERN (Conseil Européen pour la Recherche Nucléaire) in Geneva produced 290,000 bubble chamber photographs of which roughly 100 were considered to provide evidence for the existence of neutral currents. The notable point in this story is that there is no part of the model (provided by quantum field theory) that could be claimed to be isomorphic to these photographs. Weak neutral currents are the phenomenon under investigation; the photographs taken at CERN are the raw data, and any summary one might construct of the content of these photographs would be a data model. But it’s weak neutral currents that occur in the model; not any sort of data we gather in an experiment.

This is not to say that these data have nothing to do with the model. The model posits a certain number of particles and informs us about the way in which they interact both with each other and with their environment. Using this knowledge we can place them in a certain experimental context. The data we then gather in an experiment are the product of the elements of the model and of the way in which they operate in that context. Characteristically this context is one which we are able to control and about which we have reliable knowledge (e.g. knowledge about detectors, accelerators, photographic plates and so on). Using this and the model we can derive predictions about what the outcomes of an experiment will be. But, and this is the salient point, these predictions involve the entire experimental set-up and not only the model and there is nothing in the model itself with which one could compare the data. Hence, data are highly contextual and there is a big gap between observable outcomes of experiments and anything one might call a substructure of a model of neutral currents.

To underwrite this claim Bogen and Woodward notice that parallel to the research at CERN, the NAL (National Accelerator Laboratory) in Chicago also performed an experiment to detect weak neutral currents, but the data obtained in that experiment were quite different. They consisted of records of patterns of discharge in electronic particle detectors. Though the experiments at CERN and at NAL were totally different and as a consequence the data gathered had nothing in common, they were meant to provide evidence for the same theoretical model. But the model, to reiterate the point, does not contain any of these contextual factors. It posits certain particles and their interaction with other particles, not how detectors work or what readings they show. That is, the model is not idiosyncratic to a special experimental context in the way the data are and therefore it is not surprising that they do not contain a substructure that is isomorphic to the data. For this reason, models represent phenomena, not data.

It is difficult to give a general characterisation of phenomena because they do not belong to one of the traditional ontological categories (Bogen and Woodward 1988, p. 321). In fact, phenomena fall into many different established categories, including particular objects, features, events, processes, states, states of affairs, or they defy classification in these terms altogether. This, however, does not detract from the usefulness of the concept of a phenomenon because specifying one particular

ontological category to which all phenomena belong is inessential to the purpose of this section. What matters to the problem at hand is the distinctive role they play in connection with representation.

What then is the significance of data, if they are not the kind of things that models represent? The answer to this question is that data perform an evidential function. That is, data play the role of evidence for the presence of certain phenomena. The fact that we find a certain pattern in a bubble chamber photograph is evidence for the existence of neutral currents. Thus construed, we do not denigrate the importance of data in science, but we do not have to require that data have to be embeddable into the model at stake.

Those who want to establish data models as targets can reply to this in three ways. The first reply is an appeal to radical empiricism. By postulating phenomena over and above data we leave the firm ground of observable things and started engaging in trans-empirical speculation. But science has to restrict its claims to observables and remain silent (or at least agnostic) about the rest. Therefore, so the objection goes, phenomena are chimeras that cannot be part of any serious account of science. It is, however, doubtful that this helps the data model theorist. Firstly, note that it even rules out representing ‘observable phenomena’. To borrow van Fraassen’s example, on this story, a population model of deer reproduction would represent data, rather than deer (2008, pp. 254-260). Traditionally, empiricists would readily accept that deer, and the rates at which they reproduce, are observable phenomena. Denying that they are represented, by replacing them with data models, seems to be an implausible move. Secondly, irrespective of whether one understands phenomena realistically (Bogen and Woodward 1988) or antirealistically (McAllister 1997), it is phenomena that models portray and not data. To deny the reality of phenomena just won’t make a theoretical model *represent* data. Whether we regard neutral currents as real or not, it is neutral currents that are portrayed in a field-theoretical model, not bubble chamber photographs. Of course, one can suspend belief about the reality of these currents, but that is a different matter.

The second reply is to invoke a chain of representational relationships. Brading and Landry (2006) point out that the connection between a model and the world can be

broken down in two parts: the connection between a model and a data model, and the connection between a data model and the world (2006, p. 575). So the structuralist could claim that scientific models represent data models in virtue of an isomorphism between the two and additionally claim that data models in turn represent phenomena. But the key questions that need to be addressed here are (a) what establishes the representational relationship between data models and phenomena, and (b) why if a scientific model represented some data model, which in turn represented some phenomenon, would that establish a representational relationship between the model and the phenomenon itself. With respect to the first question, Brading and Landry argue that it cannot be captured within the structuralist framework (2006, p. 575). The question has just been pushed back: rather than asking how a scientific model *qua* mathematical structure represents a phenomenon, we now ask how a data model *qua* mathematical structure represents a phenomenon. With respect to the second question, although representation is not intransitive, it is not transitive (Frigg 2002, pp. 11-12). So more needs to be said regarding how a scientific model representing a data model, which in turn represents the phenomenon from which data are gathered, establishes a representational relationship between the first and last element in the representational chain.

The third reply is due to van Fraassen (2008). His ‘Wittgensteinian’ solution is to diffuse the loss of reality objection. Once we pay sufficient attention to the pragmatic features of the contexts in which scientific and data models are used, van Fraassen claims, there actually is no difference between representing data and representing a target (or a phenomenon in Bogen and Woodward’s sense): ‘in a context in which a given [data] model is *someone’s* representation of a phenomenon, there is *for that person* no difference between the question *whether a theory* [theoretical model] *fits that representation* and the question *whether that theory fits the phenomenon.*’ (2008, p. 259) Van Fraassen’s argument for this claim is long and difficult and we cannot fully investigate it here; we restrict attention to one crucial ingredient and refer the reader to Nguyen’s (2016) for a detailed discussion of the argument.

Moore’s paradox is that we cannot assert sentences of the form ‘*p* and I don’t believe that *p*’, where *p* is an arbitrary proposition. For instance, someone cannot assert that Napoleon was defeated in the battle of Waterloo and assert, at the same time, that she

doesn't believe that Napoleon was defeated in the battle of Waterloo. Van Fraassen's treatment of Moore's paradox is that speakers cannot assert such sentences because the pragmatic commitments incurred by asserting the first conjunct include that the speaker believe that p . This commitment is then contradicted by the assertion of the second conjunct. So instances of Moore's paradox are pragmatic contradictions. Van Fraassen then draws an analogy between this paradox and the scientific representation. He submits that a user simply cannot, on pain of pragmatic contradiction, assert that a data model of a target system be embeddable within a theoretical model without thereby accepting that the theoretical model represents the target.

However, Nguyen (2016) argues that in the case of using a data model as a representation of a phenomenon, no such pragmatic commitment is incurred, and therefore no such contradiction follows when accompanied by doubt that the theoretical model also represents the phenomenon. To see why this is the case, consider a more mundane example of representation: a caricaturist can represent Margaret Thatcher as draconian without thereby committing himself to the belief that Margaret Thatcher really is draconian. Pragmatically speaking, acts of representation are weaker than acts of assertion: they do not incur the doxastic commitments required for van Fraassen's analogy to go through. So it seems van Fraassen doesn't succeed in dispelling the loss of reality objection. How target systems enter the picture in the structuralist account of scientific representation remains therefore a question that structuralists who invoke data models as providing the target end structures must address. Without such an account the structuralist account of representation remains at the level of data, a position that seems implausible, and contrary to actual scientific practice.

We now turn to the second response: that a structure is instantiated in the system. As mentioned above, this response comes in three versions. The first is metaphysically more parsimonious and builds on the systems' constituents. Although target systems are not structures, they are composed of parts that instantiate physical properties and relations. The parts can be used to define the domain of individuals, and by considering the physical properties and relations purely extensionally, we arrive at a class of extensional relations defined over that domain (see for instance Suppes'

discussion of the solar system (2002, p. 22)). This supplies the required notion of structure. We might then say that physical systems instantiate a certain structure, and it is this structure that models are isomorphic to.

As an example consider the methane molecule. The molecule consists of a carbon atom and four hydrogen atoms grouped around it, forming a tetrahedron. Between each hydrogen atom and the carbon atom there is a covalent bond. One can then regard the atoms as objects and the bonds are relations. Denoting the carbon atom by a , and the four hydrogen atoms by b , c , d , and e , we obtain a structure S with the domain $U = \{a, b, c, d, e\}$ and the relation $r = \{ \langle a, b \rangle, \langle b, a \rangle, \langle a, c \rangle, \langle c, a \rangle, \langle a, d \rangle, \langle d, a \rangle, \langle a, e \rangle, \langle e, a \rangle \}$, which can be interpreted as ‘being connected by a covalent bond’.

The main problem facing this approach is the underdetermination of target-end structure. Underdetermination threatens in two distinct ways. Firstly, in order to identify the structure determined by a target system, a domain of objects is required. What counts as an object in a given target system is a substantial question (Frigg 2006). One could just as well choose bonds as objects and consider the relation ‘sharing a node with another bond’. Denoting the bonds by a' , b' , c' and d' , we obtain a structure S' with the domain $U' = \{a', b', c', d'\}$ and the relation $r = \{ \langle a', b' \rangle, \langle b', a' \rangle, \langle a', c' \rangle, \langle c', a' \rangle, \langle a', d' \rangle, \langle d', a' \rangle, \langle b', c' \rangle, \langle c', b' \rangle, \langle b', d' \rangle, \langle d', b' \rangle, \langle c', d' \rangle, \langle d', c' \rangle \}$. Obviously S and S' are not isomorphic. So which structure is picked out depends on how the system is described. Depending on which parts one regards as individuals and what relation one chooses, very different structures can emerge. And it takes little ingenuity to come up with further descriptions of the methane molecule, which lead to yet other structures.

There is nothing special about the methane molecule, and any target system can be presented under alternative descriptions, which ground different structures. So the lesson learned generalises: there is no such thing as *the* structure of a target system. Systems only have a structure under a particular description, and there are many non-equivalent descriptions. This renders talk about a model being isomorphic to target system *simpliciter* meaningless. Structural claims do not ‘stand on their own’ in that

their truth rests on the truth of a more concrete description of the target system. As a consequence, descriptions are an integral part of an analysis of scientific representation.

In passing we note that Frigg (2006, pp. 55-56) also provides another argument that pulls into the same direction: structural claims are abstract and are true only relative to a more concrete non-structural description. For a critical discussion of this argument see Frisch's (2015, pp. 289-294) and Portides' (forthcoming).

How much of a problem this is depends on how austere one's conception of models is. The semantic view of theories was in many ways the result of an anti-linguistic turn in the philosophy of science. Many proponents of the view aimed to exorcise language from an analysis of theories, and they emphasised that the model-world relationship ought to be understood as a *purely* structural relation. Van Fraassen, for instance, submits that 'no concept which is essentially language dependent has any philosophical importance at all' (1980, p. 56) and observes that '[t]he semantic view of theories makes language largely irrelevant' (1989, p. 222). And other proponents of the view, while less vocal about the irrelevance of language, have not assigned language a systematic place in their analysis of theories.

For someone of that provenance the above argument is bad news. However, a more attenuated position could integrate descriptions in the package of modelling, but this would involve abandoning the idea that representation can be cashed out solely in structural terms. Bueno and French have recently endorsed such a position. They accept the point that different descriptions lead to different structures and explain that such descriptions would involve 'at the very least some minimal mathematics and certain physical assumptions' (Bueno and French 2011, p. 887). Likewise, 'Munich' structuralists explicitly acknowledge the need for a concrete description of the target-system (Balzer et al. 1987, pp. 37-38), and they consider these 'informal descriptions' to be 'internal' to the theory. This is a plausible move, but those endorsing this solution have to concede that there is more to epistemic representation than structures and morphisms.

The second way in which structural indeterminacy can surface is via Newman's theorem. The theorem essentially says that any system instantiates any structure, the only constraint being cardinality (a practically identical conclusion is reached in Putnam's so called model-theoretic argument; see Demopoulos' (2003) for a discussion). Hence, *any* structure of cardinality C is isomorphic to a target of cardinality C because the target instantiates any structure of cardinality C (see Ketland's (2004) and Frigg and Votsis' (Frigg and Votsis 2011) for discussions). This problem is not unsolvable, but all solutions require that among all structures formally instantiated by a target system one is singled out as being the true or natural structure of the system. How to do this in the structuralist tradition remains unclear (Ainsworth (2009) provides a useful summary of the different solutions).

Newman's theorem is both stronger and weaker than the argument from multiple descriptions. It's stronger in that it provides more alternative structures than multiple descriptions. It's weaker in that many of the structures it provides are 'unphysical' because they are purely set theoretical combinations of elements. By contrast, descriptions pick out structures that a system can reasonably be seen as possessing.

The second version of the second response emerges from the literature on the applicability of mathematics. Structural Platonists like Resnik (1997) and Shapiro (1983; 1997; 2000) take structures to be 'ante rem' universals. On this view, structures exist independently of physical systems, yet they can be instantiated in physical systems. On this view systems instantiate structures and models are isomorphic to these instantiated structures.

This view raises all kind of metaphysical issues about the ontology of structures and the instantiation relation. Let us set aside these issues and assume that they can be resolved in one way or another. This would still leave us with serious epistemic and semantic questions. How do we know a certain structure is instantiated in a system and how do we refer to it? Objects do not come with labels on their sleeves specifying which structures they instantiate, and proponents of structural universals face a serious problem in providing an account of *how we access* the structures instantiated by target systems. Even if – as a brute metaphysical fact – target systems only instantiate a small number of structures, and therefore there is a substantial question

regarding whether or not scientific models represent them, this does not help us understand how we could ever come to know whether or not the isomorphism holds. It seems that individuating a domain of objects and identifying relations between them is the only way for us to access a structure. But then we are back to the first version of the response, and we are again faced with all the problems that it raises.

The third version of the second response is more radical. One might take target systems themselves to be structures. If this is the case then there is no problem with the idea that they can be isomorphic to a scientific model. One might expect ontic structural realists to take this position. If the world fundamentally is a structure, then there is nothing mysterious about the notion of an isomorphism between a model and the world. Surprisingly, some ontic structuralists have been hesitant to adopt such a view (see French and Ladyman (1999, p. 113) and French (2014, p. 195)). Others, however, seem to endorse it. Tegmark (2008), for instance, offers an explicit defence of the idea that the world simply is a mathematical structure. He defines a seemingly moderate form of realism – what he calls the ‘external reality hypothesis (ERH)’ – as the claim that ‘there exists an external physical reality completely independent of us humans’ (*ibid.*, p. 102) and argues that this entails that the world is a mathematical structure (his ‘mathematical universe hypothesis’) (*ibid.*, p. 102). His argument for this is based on the idea that a so called ‘theory-of-everything’ must be expressible in a form that is devoid of human-centric ‘baggage’ (by the ERH), and the only theories that are devoid of such baggage are mathematical, which, strictly speaking, describe mathematical structures. Thus, since a complete theory of everything describes an external reality independent of humans, and since it describes a mathematical structure, the external reality itself *is* a mathematical structure.

This approach stands or falls on the strengths of its premise that a complete theory of everything will be formulated purely mathematically, without any ‘human baggage’, which in turn relies on a strict reductionist account of scientific knowledge (*ibid.*, pp. 103-104). Discussing this in any detail goes beyond our current purposes. But it is worth noting that Tegmark’s discussion is focused on the claim that *fundamentally* the world is a mathematical structure. Even if this were the case, it seems irrelevant for many of our current scientific models, whose targets aren’t at this level. When modelling an aeroplane wing we don’t refer to the fundamental super-string structure

of the bits of matter that make up the wing, and we don't construct wing models that are isomorphic to such fundamental structures. So Tegmark's account offers no answer to the question about where structures are to be found at the level of non-fundamental target systems.

6. The Inferential Conception

In this section we discuss accounts of scientific representation that analyse representation in terms of the inferential role of scientific models. On the previous accounts discussed, a model's inferential capacity dropped out of whatever it was that was supposed to answer the ER-problem: proposed morphisms or similarity relations between models and their targets for example. The accounts discussed in this section build the notion of surrogative reasoning directly into the conditions on epistemic representation.

6.1 Deflationary Inferentialism

Suárez argues that we should adopt a 'deflationary or minimalist attitude and strategy' (2004, p. 770) when addressing the problem of epistemic representation. We will discuss deflationism in some detail below, but in order to formulate and discuss Suárez's theory of representation we need at least a preliminary idea of what is meant by a deflationary attitude. In fact two different notions of deflationism are in operation in his account. The first is 'abandoning the aim of a substantive theory to seek universal necessary and sufficient conditions that are met in each and every concrete real instance of scientific representation ... necessary conditions will certainly be good enough.' (ibid., p. 771) We call the view that a theory of representation should provide only necessary conditions n-deflationism ('n' for 'necessary'). The second notion is that we should seek 'no deeper features to representation other than its surface features' (ibid., p. 771) or 'platitudes' (Suárez and Solé 2006, p. 40), and that we should deny that an analysis of a concept 'is the kind of analysis that will shed explanatory light on our use of the concept' (Suárez 2015, p. 39). We call this position s-deflationism ('s' for 'surface feature'). As far as

we can tell, Suárez intends his account of representation to be deflationary in both senses.

Suárez dubs the account that satisfies these criteria ‘inferentialism’ (2004, p. 773):

Inferentialism 1: A scientific model M represents a target T only if (i) the representational force of M points towards T , and (ii) M allows competent and informed agents to draw specific inferences regarding T .

Notice that this condition is not an instantiation of the ER-Scheme: in keeping with n-deflationism it features a material conditional rather than a biconditional and hence provides necessary (but not sufficient) conditions for M to represent T . We now discuss each condition in turn, trying to explicate in what way they satisfy s-deflationism.

The first condition is designed to make sure that M and T indeed enter into a representational relationship, and Suárez stresses that representational force is ‘necessary for any kind of representation’ (ibid., p. 776). But explaining representation in terms of representational force seems to shed little light on the matter as long as no analysis of representational force is offered. Suarez addresses this point by submitting that the first condition can be ‘satisfied by mere stipulation of a target for any source’ (ibid., p. 771). This might look like denotation as in Section 3. But Suárez stresses that this is not what he intends for two reasons. Firstly, he takes denotation to be a substantive relation between a model and its target, and the introduction of such a relation would violate the requirement of s-deflationism (Suárez 2015, p. 41). Secondly, M can denote T only if T exists. Thus including denotation as a necessary condition on scientific representation ‘would rule out fictional representation, that is, representation of nonexisting entities’ (Suárez 2004, p. 772), and ‘any adequate account of scientific representation must accommodate representations with fictional or imaginary targets’ (Suárez 2015, p. 44).

The second issue is one that besets other accounts of representation too, in particular similarity and isomorphism accounts. The first reason, however, goes right to the heart of Suárez’s account: it makes good on the s-deflationary condition that nothing

other than surface features can be included in an account of representation. At a surface level one cannot explicate ‘representational force’ at all and any attempt to specify what representational force consists in is a violation of s-deflationism.

The second necessary condition, that models allow competent and informed agents to draw specific inferences about their targets, is in fact just the *Surrogate Reasoning Condition* we introduced in Section 2, now taken as a necessary condition on epistemic representation. The sorts of inferences that models allow are not constrained. Suárez points out that the condition ‘does not require that [M] allow deductive reasoning and inference; any type of reasoning inductive, analogical, abductive – is in principle allowed’ (Suárez 2004, p. 773). (The insistence on inference makes Suárez’s account an instance of what Chakravartty (2010) calls a ‘functional conception’ of representation.)

A problem for this approach is that we are left with no account of how these inferential rules are generated: what is it about models that allows them to licence inferences about their targets, or what leads them to licence some inferences and not others? Contessa makes this point most stridently when he argues that:

‘On the inferential conception, the user’s ability to perform inferences from a vehicle [model] to a target seems to be a brute fact, which has no deeper explanation. This makes the connection between epistemic representation and valid surrogate reasoning needlessly obscure and the performance of valid surrogate inferences an activity as mysterious and unfathomable as soothsaying or divination’ (Contessa 2007, p. 61)

This seems correct, but Suárez can dismiss this complaint by appeal to s-deflationism. Since inferential capacity is supposed to be a surface level feature of scientific representation, we are not supposed to ask for any elucidation about what makes an agent competent and well informed and how inferences are drawn.

For these reasons Suárez’s account is deflationary both in the sense of n-deflationism and of s-deflationism. His position provides us with a concept of epistemic representation that is cashed out in terms of an inexplicable notion of representational force and of an inexplicable capacity to ground inferences. This is very little indeed. It is the adoption of a deflationary attitude that allows him to block any attempt to

further unpack these conditions and so the crucial question is: why should one adopt deflationism?

We turn to this question shortly. Before doing so we want to briefly outline how the above account fares with respect to the other problems introduced in Section 2. The account provides a neat explanation of the possibility of misrepresentation: ‘Part (ii) of this conception accounts for inaccuracy since it demands that we correctly draw inferences from the source about the target, but it does not demand that the conclusions of these inferences be all true, nor that all truths about the target may be inferred’ (Suárez 2004, p. 776). Models represent their targets only if they license inferences about them. They represent them accurately to the extent that the conclusions of these inferences are true.

With respect to the representational demarcation problem, Suárez illustrates his account with a large range of representations, including diagrams, equations, scientific models, and non-scientific representations such as artistic portraits. He explicitly states that ‘if the inferential conception is right, scientific representation is in several respects very close to iconic modes of representation like painting’ (Suárez 2004, p. 777) and he mentions the example of Velázquez’s portrait of Innocent X (ibid.) It is clear that the conditions of *Inferentialism 1* are met by non-scientific as well as scientific epistemic representations. So, at least without sufficient conditions, there is no clear way of demarcating between the different kinds of epistemic representation.

Given the wide variety of types of representation that this account applies to, it’s unsurprising that Suarez has little to say about the ontological problem. The only constraint that *Inferentialism 1* places on the ontology of models is that ‘[i]t requires $[M]$ to have the internal structure that allows informed agents to correctly draw inferences about $[T]$ ’ (Suárez 2004, p. 774). And relatedly, since the account is supposed to apply to a wide variety of entities, including equations and mathematical structures, the account implies that mathematics is successfully applied in the sciences, but in keeping with the spirit of deflationism no explanation is offered about how this is possible.

Suárez does not directly address the problem of style, but a minimalist answer emerges from what he says about representation. On the one hand he explicitly acknowledges that many different kinds of inferences are allowed by the second condition in *Inferentialism 1*. In the passage quoted above he mentions inductive, analogical and abductive inferences. This could be interpreted as the beginning of classification of representational styles. On the other hand, Suárez remains silent about what these kinds are and about how they can be analysed. This is unsurprising because spelling out what these inferences are, and what features of the model ground them, would amount to giving a substantial account, which is something Suárez wants to avoid.

Let us now return to the question about the motivation for deflationism. As we have seen, a commitment to deflationism about the concept is central to Suárez's approach to scientific representation. But deflationism comes in different guises, which Suárez illustrates by analogy with deflationism with respect to truth. Suárez (2015) distinguishes between the 'redundancy' theory (associated with Frank Ramsey and also referred to as the 'no theory' view), 'abstract minimalism' (associated with Crispin Wright) and the 'use theory' (associated with Paul Horwich). What all three are claimed to have in common is that they accept the disquotational schema – i.e. instances of the form: ' P ' is true iff P . Moreover they 'either do not provide an analysis in terms of necessary and sufficient conditions, or if they do provide such conditions, they claim them to have no explanatory purchase' (ibid., p. 37).

He claims that the redundancy theory of truth is characterised by the idea that 'the terms 'truth' and 'falsity' do not admit a theoretical elucidation or analysis. But that, since they can be eliminated in principle – if not in practice – by disquotation, they do not in fact require such an analysis' (ibid., p. 39). So, as Suarez characterises the position, the redundancy theory denies that any necessary and sufficient conditions for application of the truth predicate case be given. He argues that: 'the generalization of this 'no-theory theory' for any given putative concept X is the thought that X neither possesses nor requires necessary and sufficient conditions because it is not in fact a 'genuine', explanatory or substantive concept' (ibid). This motivates n-deflationism. (Although one might ask why such a position would allow even necessary conditions. Suárez doesn't discuss this.)

This approach faces a number of challenges. First, the argument is based on the premise that if deflationism is good for truth it must be good for representation. This premise is assumed tacitly. There is, however, a question whether the analogy between truth and representation is sufficiently robust to justify subjecting them to the same theoretical treatment. Surprisingly, Suárez offers little by way of explicit argument in favour of any sort of deflationary account of epistemic representation. In fact, the natural analogue of the linguistic notion of truth is accurate epistemic representation, rather than epistemic representation itself, which may be more appropriately compared with linguistic meaning. Second, the argument insinuates that deflationism is the correct analysis of truth. This, however, is far from an established fact. Different positions are available in the debate and whether deflationism (or any specific version of it) is superior to other proposals remains a matter of controversy (see, for instance, Küne's (2003)). But as long as it's not clear that deflationism about truth is a superior position, it's hard to see how one can muster support for deflationism about representation by appealing to deflationism about truth.

Moreover, a position that allows only necessary conditions on epistemic representation faces a serious problem. While such an account allows us to *rule out* certain scenarios as instances of epistemic representation (for example a proper name doesn't allow for a competent and well informed language user to draw any specific inferences about its bearer and Callender and Cohen's salt-shaker doesn't allow a user to draw any specific inferences about Madagascar), the lack of sufficient conditions doesn't allow us to *rule in* any scenario as an instance of epistemic representation. So on the basis of *Inferentialism I* we are never in position to assert that a particular model actually is a representation, which is an unsatisfactory situation.

The other two deflationary positions in the debate over truth are abstract minimalism and the use theory. Suárez characterises the use theory as being based on the idea that 'truth is nominally a property, although not a substantive or explanatory one, which is essentially defined by the platitudes of its use of the predicate in practice (2015, p. 40). Abstract minimalism is presented as the view that while truth is 'legitimately a property, which is abstractly characterized by the platitudes, it is a property that cannot explain anything, in particular it fails to explain the norms that govern its very

use in practice' (ibid., p. 40). Both positions imply that necessary and sufficient conditions for truth *can* be given (ibid.). But on either account, such conditions only capture non-explanatory surface features. This motivates s-deflationism.

Since s-deflationism explicitly allows for necessary and sufficient conditions, *Inferentialism 1* can be extended to an instance of the ER-scheme, providing necessary and sufficient conditions (which also seems to be in line with Suárez and Solé (2006, p. 41) who provide a formulation of inferentialism with a biconditional):

Inferentialism 2: A scientific model M represents a target T iff (i) the representational force of M points towards T , and (ii) M allows competent and informed agents to draw specific inferences regarding T .

If one takes conditions (i) and (ii) to refer to 'features of activates within a normative practise, [that] do not stand for relations between sources and targets' (Suárez 2015, p. 46), then we arrive at a 'use-based' account of epistemic representation. In order to understand a particular instance of a model M representing a target T we have to understand how scientists go about establishing that M 's representational force points towards T , and the inferential rules, and particular inferences from M to T , they use and make.

Plausibly, such a focus on practice amounts to looking at the inferential rules employed in each instance, or type of instance, of epistemic representation. This, however, raises a question about the status of any such analysis vis-à-vis the general theory of representation as given in *Inferentialism 2*. There seem to be two options. The first is to affirm *Inferentialism 2*'s status as an exhaustive theory of representation. This, however, would imply that any analysis of the workings of a particular model would fall outside the scope of a theory of representation because any attempt to address Contessa's objection would push the investigation outside the territory delineated by s-deflationism. Such an approach seems to be overly purist. The second option is to understand *Inferentialism 2* as providing abstract conditions that require concretization in each instance of epistemic representation (abstraction can here be understood, for instance, in Cartwright's (1999) sense). Studying the concrete realisations of the abstract conditions is then an integral part of the theory.

This approach seems plausible, but it renders deflationism obsolete. Thus understood, the view becomes indistinguishable from a theory that accepts the *Surrogate Reasoning Condition* and the *Requirement of Directionality* as conditions of adequacy and analyses them in pluralist spirit, that is, under the assumption that these conditions can have different concrete realisers in different contexts. But this programme can be carried out without ever mentioning deflationism.

One might reply that the first option unfairly stacks the deck against inferentialism and point out that different inferential practices *can* be studied within the inferentialist framework. One way of making good on this idea would be to submit that the inferences from models to their targets should be taken as conceptually basic, denying that they need to be explained; in particular, denying that they need to be grounded by any (possibly varying) relation(s) that might hold between models and their targets. Such an approach is inspired by Brandom's inferentialism in the philosophy of language where the central idea is to reverse the order of explanation from representational notions – like truth and reference – to inferential notions – such as the validity of argument (1994; 2000). Instead, we are urged to begin from the inferential role of sentences (or propositions, or concepts, and so on) – that is the role that they play in providing reasons for other sentences (or propositions etc.), and having such reasons provided for them – and from this reconstruct their representational aspects.

Such an approach is developed by de Donato Rodríguez and Zamora Bonilla (2009) and seems like a fruitful route for future research, but for want of space we will not discuss it in detail here. There is no evidence that Suarez would endorse such an approach. And, more worrying for *Inferentialism 2*, it is not clear whether such an approach would satisfy s-deflationism. Each investigation into the inferential rules utilised in each instance, or type of instance of epistemic representation will likely be a substantial (possibly sociological or anthropological) project. Thus the s-deflationary credentials of the approach – at least if they are taken to require that nothing substantial can be said about scientific representation in each instance, as well as in general – are called into question.

Finally, if the conditions in *Inferentialism 2* are taken to be abstract platitudes then we arrive at an abstract minimalism. Although *Inferentialism 2* defines the concept of epistemic representation, the definition does not suffice to explain the use of any particular instance of epistemic representation for: ‘on the abstract minimalism here considered, to apply this notion to any given concrete case of representation requires that some additional relation obtains between $[M]$ and $[T]$, or a property of $[M]$ or $[T]$, or some other application condition’ (Suárez 2015, p. 48; *cf.* Suárez and Solé 2006). Hence, according to this approach representational force and inferential capacity are taken to be abstract platitudes that suffice to define the concept of scientific representation. However, because of their level of generality, they fail to explain any particular instance of it. To do this requires reference to additional features that vary from case to case. These other conditions can be ‘isomorphism or similarity’ and they ‘would need to obtain in each concrete case of representation’ (Suárez 2015, p. 45) (*cf.* (Suárez 2004, p. 776), (Suárez 2004, p. 773) and (Suárez and Solé 2006, p. 43)). These extra conditions are called the *means* of representation, the relations that scientists exploit in order to draw inferences about targets from their models, and are to be distinguished from conditions (i) and (ii), the *constituents* of representation, that define the concept (Suárez 2003, p. 230; 2010, pp. 93-94; 2015, p. 46; Suárez and Solé 2006, p. 43). We are told that the means cannot be reduced to the constituents but that ‘all representational means (such as isomorphism and similarity) are concrete instantiations, or realisations, of one of the basic platitudes that constitute representation’ (Suárez and Solé 2006, p. 43) and that ‘there can be no application of representation without the simultaneous instantiation of a particular set of properties of $[M]$ and $[T]$, and their relation’ (Suárez and Solé 2006, p. 44).

Such an approach amounts to using conditions (i) and (ii) to answer the ER-problem, but again with the caveat that they are abstract conditions that require concretisation in each instance of epistemic representation. In this sense it is immune to Contessa’s objection about the ‘mysterious’ capacity that models have to licence about their targets. They do so in virtue of more concrete relations that hold between models and their targets, albeit relations that vary from case to case. The key question facing this account is to fill in the details about what sort of relations concretise the abstract conditions. But we are now facing a similar problem as the above. Even if s-deflationism applies to epistemic representation in general, an investigation into each

specific instance of will involve uncovering substantial relations that hold between models and their targets, which again conflicts with Suárez's adherence to deflationist approach.

6.2 Inflating Inferentialism: Interpretation

In response to difficulties like the above Contessa claims that 'it is not clear why we should adopt a deflationary attitude *from the start*' (2007, p. 50) and provides a 'interpretational account' of scientific representation that is still, at least to some extent, inspired by Suárez's account, but without being deflationary. Contessa claims:

'[t]he main difference between the interpretational conception [...] and Suárez's inferential conception is that the interpretational account is a substantial account — interpretation is not just a "symptom" of representation; it is what makes something an epistemic representation of a something else.' (ibid., p. 48)

To explain in virtue of what the inferences can be drawn, Contessa introduces the notion of an *interpretation* of a model, in terms of its target system as a necessary and sufficient condition on epistemic representation:

Interpretation: 'A scientific model M is an epistemic representation of a certain target T (for a certain user) if and only if the user adopts an interpretation of M in terms of T .' (Contessa 2007, p. 57; see also Contessa 2011, pp. 126-127)

Contessa offers a detailed formal characterisation of an interpretation, which we cannot repeat here for want of space (see (Contessa 2007, pp. 57-62) for details). The leading idea is that the model user first identifies a set of relevant objects in the model, and a set of properties and relations these objects instantiate, along with a set of relevant objects in the target and a set of properties and relations these objects instantiate. The user then (a) takes M to denote T ; (b) takes every identified object in the model to denote exactly one object in the target (and every relevant object in the target has to be so denoted and as a result there is a one-to-one correspondence between relevant objects in the model and relevant objects in the target); (c) takes every property and relation in the model to denote a property or relation of the same

arity in the target (and, again, and every property and relation in the target has to be so denoted and as a result there will be one-to-one correspondence between relevant properties and relations in the model and target). A formal rendering of these conditions is what Contessa calls an ‘analytic interpretation’ (he also includes an additional condition pertaining to functions in the model and target, which we suppress for brevity). The relationship between interpretations and the surrogative reasoning mentioned above is that it is *in virtue* of the user adopting an analytic interpretation that a model licences inferences about its target.

At first sight Contessa’s interpretation may appear to be equivalent to setting up an isomorphism between model and target. This impression is correct in as far as an interpretation requires that there be a one-to-one correspondence between relevant elements and relations in the model and the target. However, unlike the isomorphism view, Contessa’s interpretations are not committed to models being structures, and relations can be interpreted full fledged relations rather than purely extensionally specified sets of tuples.

Interpretation is a non-deflationary account of scientific representation: most (if not all) instances of scientific representation involve a model user adopting an analytic interpretation towards a target. The capacity for surrogative reasoning is then seen as a symptom of the more fundamental notion of a model user adopting an interpretation of a model in terms of its target. For this reason the adoption of an analytical interpretation is a substantial sufficient condition on establishing the representational relationship. Contessa focuses on the sufficiency of analytic interpretations rather than their necessity and adds that he does ‘not mean to imply that all interpretation of vehicles [models] in terms of the target are necessarily analytic. Epistemic representations whose standard interpretations are not analytic are at least conceivable’ (Contessa 2007, p. 58). Even with this in mind, it is clear that he intends that there be *some* interpretation is a necessary condition on epistemic representation.

Let’s now turn to how *Interpretation* fares with respect to our questions for an account of epistemic representation as set out in section 2. Modulo the caveat about non-analytical interpretations, *Interpretation* provides necessary and sufficient conditions on epistemic representation and hence answers the ER-Problem. Furthermore, it does

so in a way that explains the directionality of representation: interpreting a model in terms of a target does not entail interpreting a target in terms of a model.

Contessa does not comment on the applicability of mathematics but since his account shares with the structuralist account an emphasis on relations and one-to-one model-target correspondence, Contessa can appeal to the same account of the applicability of mathematics as structuralist.

With respect to the demarcation problem, Contessa is explicit that ‘[p]ortraits, photographs, maps, graphs, and a large number of other representational devices’ perform inferential functions (Contessa 2007, p. 54). Since nothing in the notion of an interpretation seems restricted to scientific models, it is plausible to regard *Interpretation* as a universal theory of epistemic representation (a conclusion that is also supported by the fact that Contessa (ibid.) uses the example of the London underground map to motivate his account; see also (Contessa 2011)). As such, *Interpretation* seems to deny the existence of a substantial distinction between scientific and non-scientific epistemic representations (at least in terms of their representational properties). It remains unclear how *Interpretation* addresses the problem of style. As we have seen earlier, in particular visual representations fall into different categories. It is a question for future research how these can be classified within the interpretational framework.

With respect to the question of ontology, *Interpretation* itself places few constraints on what scientific models are, ontologically speaking. All it requires is that they consist of objects, properties, relations, and functions. For this reason our discussion in subsection 4.3 above rears its head again here. As before, how to apply *Interpretation* to physical models can be understood relatively easily. But how to apply it to non-physical models is less straightforward. Contessa (2010) distinguishes between mathematical models and fictional models, where fictional models are taken to be fictional objects. We briefly return to his ontological views in Section 7.

In order to deal with the possibility of misrepresentation, Contessa notes that ‘a user does not need to believe that every object in the model denotes some object in the system in order to interpret the model in terms of the system’ (2007, p. 59). He

illustrates this claim with an example of contemporary scientists using the Aristotelian model of the cosmos to represent the universe, pointing out that ‘in order to interpret the model in terms of the universe, we do not need to assume that the sphere of fixed stars itself [...] denotes anything in the universe’ (ibid.).

From this example it is clear that the relevant sets of objects, properties and functions isolated in the construction of the analytic interpretation do not need to exhaust the objects, properties, relations, and functions of either the model or the target. The model user can identify a relevant *proper* subset in each instance. This allows *Interpretation* to capture the common practise of abstraction in scientific models: a model need only represent some features of its target, and moreover, the model may have the sort of ‘surplus’ features are not taken to represent anything in the target, i.e. that not all of a model’s features need to play a direct representational role.

This suggestion bears some resemblance to partial structures, and it suffers from the same problem too. In particular distortive idealisations are a source of problems for *Interpretation*, as several commentators have observed (*cf.* Shech’s (2014) and Bolinska’s (2013)). Contessa is aware of this problem and illustrates it with the example of a massless string. His response to the problem is to appeal to a user’s corrective abilities: ‘Since models often misrepresent some aspect of the system or other, it is usually up to the user’s competence, judgment, and background knowledge to use the model successfully in spite of the fact that the model misrepresents certain aspects of the system.’ (Contessa 2007, p. 60) This is undoubtedly true, but it is unclear how such a view relates, or even derives from, *Interpretation*. An appeal to the competence of users seems to be an *ad hoc* move that has no systematic grounding in the idea of an interpretation, and it is an open question how the notion of an interpretation could be amended to give distortive idealisations a systematic place.

Ducheyne (2012) provides a variant of *Interpretation* that one might think could be used to accommodate these distortive idealisations. The details of the account, which we won’t state precisely here for want of space, can be found in (Ducheyne 2012, pp. 83-86). The central idea is that each relevant relation specified in the interpretation, holds precisely in the model, and corresponds to the same relation that holds only approximately (with respect to a given purpose) in the target. For example, the low

mass of an actual pendulum's string approximates the masslessness of the string in the model. The one-to-one correspondence between (relevant) objects and relations in the model and target is retained, but the notion of a user taking relations in the model to denote relations in the target, is replaced with the idea that the relations in the target are approximations of the ones they correspond to. Ducheyne calls this the *Pragmatic Limiting Case* account of scientific representation (the pragmatic element comes from the fact that the level of approximation required is determined by the purpose of the model user).

However, if this account is to succeed in explaining how distortive idealisations are scientific representations, then more needs to be said about how a target relation can 'approximate' a model relation. Ducheyne implicitly relies on the fact that relations are such that 'we can determine *the extent to which* [they hold] empirically' (2012, p. 83, emphasis added). This suggests that he has quantifiable relations in mind, and that what it means for a relation r in the target to approximate a relation r' in the model is a matter of comparing numerical values, where a model user's purpose determines how close they must be if the former is to count as an approximation of the latter. But whether this exhausts the ways in which relations can be approximations remains unclear. Hendry (1998), Laymon (1990), Liu (1999), Norton (2012), and Ramsey (2006), among others, offer discussions of different kinds of idealisations and approximations, and Ducheyne would have to make it plausible that all these can be accommodated in his account.

More importantly, Ducheyne's account has problems dealing with misrepresentations. Although it is designed to capture models that misrepresent by being approximations, of their targets, it remains unclear how it deals with models that are outright mistaken. For example, it seems a stretch to say that Thomson's model of the atom (now derogatively referred to as the 'plum pudding model') is an approximation of what the quantum mechanical shell model tells us about atoms, and it seems unlikely that there is a useful sense in which the relations that hold between electrons in Thomson's model 'approximate' those that hold in reality. But this does not mean that it is not a scientific representation of the atom; it's just an incorrect one. It does not seem to be the case that all cases of scientific misrepresentation are instances where the model is an approximation of the target (or even conversely, it is not clear whether all instances

of approximation need to be considered cases of ‘misrepresentation’ in the sense that they licence falsehoods about their targets).

6.3 The DDI Account

Our final account is Hughes’ *Denotation, Demonstration, and Interpretation* (DDI) account of scientific representation (1997; 2010, Ch. 5). This account has inspired both the inferential (see Suárez’s (2004, p. 770) and (2015)) and the interpretational account (see Contessa’s (2011, p. 126)) discussed in this section.

Quoting directly from Goodman (1976, p. 5), Hughes takes a model of a physical system to ‘be a symbol for it, stand for it, refer to it’ (Hughes 1997, p. 330). Presumably the idea is that a model denotes its target in the same way that a proper name denotes its bearer, or, stretching the notion of denotation slightly, a predicate denotes elements in its extension. (Hughes (1997, p. 330) notes that there is an additional complication when the model has multiple targets but this is not specific to the DDI account and is discussed in more detail in Section 8.) This is the first ‘D’ in ‘DDI’. What makes models epistemic representations and thereby distinguishes them from proper names, are the demonstration and interpretation conditions.

The demonstration condition, the second ‘D’ in ‘DDI’, relies on a model being a ‘secondary subject that has, so to speak, a life of its own. In other words, [a] representation has an internal dynamic whose effects we can examine’ (1997, p. 331) (that models have an ‘internal dynamic’ is all that Hughes has to say about the problem of ontology). The two examples offered by Hughes are both models of what happens when light is passed through two nearby slits. One model is mathematical where the internal dynamics are ‘supplied by the deductive, resources of the mathematics they employ’ (ibid.), the other is a physical ripple chamber where they are supplied by ‘the natural processes involved in the propagation of water waves’ (ibid., p. 332).

Such demonstrations, on either mathematical models or physical models are still primarily about the models themselves. The final aspect of Hughes’ account – the ‘I’ in ‘DDI’ – is interpretation of what has been demonstrated in the model in terms of

the target system. This yields the predictions of the model (ibid., p. 333). Unfortunately Hughes has little to say about what it means to interpret a result of a demonstration on a model in terms of its target system, and so one has to retreat to an intuitive (and unanalysed) notion of carrying over results from models to targets.

Now Hughes is explicit that he is not attempting to answer the ER-problem, and that he does not even offer denotation, demonstration and interpretation as individually necessary and jointly sufficient conditions for scientific representation. He prefers the more 'modest suggestion that, if we examine a theoretical model with these three activities in mind, we shall achieve some insight into the kind of representation that it provides' (ibid., p. 339).

We are not sure how to interpret Hughes' position in light of this. On one reading, he can be seen as describing how we *use* models. As such, *DDI* functions as a diachronic account of what a model user does when using a model in an attempt to learn about a target system. We first stipulate that the model stands for the target, then prove what we want to know, and finally 'transfer' the results obtained in the model back to the target. Details aside, this picture seems by and large correct. The problem with the *DDI* account is that it does not explain why and how this is possible. Under what conditions is it true that the model denotes the target? What kinds of things are models that they allow for demonstrations? How does interpretation work; that is, how can results obtained in the model be transferred to the target? These are questions an account of epistemic representation has to address, but which are left unanswered by the *DDI* account thus interpreted. Accordingly, *DDI* provides an answer to a question distinct from the ER-problem. Although a valuable answer to the question of how models are used, it does not help us to much here, since it presupposes the very representational relationship we are interested in between models and their targets.

An alternative reading of Hughes' account emerges when we consider the developments of the structuralist and similarity conceptions discussed previously, and the discussion of deflationism in subsection 6.1: perhaps the very act of using a model, with all the user intentions and practices that brings with it, constitutes the epistemic representation relationship itself. And as such, perhaps the *DDI* conditions could be taken as an answer to the ER-problem:

DDI-ER: A scientific model M represents a target T iff M denotes T , an agent (or collection of thereof) S exploits the internal dynamic of M to make demonstrations D , which in turn are interpreted by the agent (or collection of thereof) to be about T .

This account comes very close to *Interpretation* as discussed in the previous subsection. And as such it serves to answer the questions we set out in Section 2 above in the same way. But in this instance, the notion of what it means to ‘exploit an internal dynamic’ and ‘interpret the results’ of this to be about T need further explication. If ‘interpretation’ is cashed out in the same way as Contessa’s analytic interpretation, then the account will be vulnerable to the same issues as those discussed previously. In another place Hughes endorses Giere’s semantic view of theories, which he characterises as connecting models to the target with a theoretical hypothesis (Hughes 1998, p. 121). This suggests that an interpretation is a theoretical hypothesis in this sense. If so, then Hughes’s account collapses into a version of Giere’s.

Given that Hughes describes his account as ‘designedly skeletal [and in need] to be supplemented on a case-by-case basis’ (Hughes 1997, p. 335), one option available is to take the demonstration and interpretation conditions to be abstract (in the sense of abstract minimalism discussed above), which require filling in each instance, or type of instance, of epistemic representation. As Hughes notes, his examples of the internal dynamics of mathematical and physical models are radically different with the demonstrations of the former utilizing mathematics, and the latter, physical properties such as the propagation of water waves. Similar remarks apply to the interpretation of these demonstrations, as well as to denotation. But as with Suárez’s account, the definition sheds little light on the problem at hand as long as no concrete realisations of the abstract conditions are discussed. Despite Hughes’ claims to the contrary, such an account could prove a viable answer the ER-problem, and it seems to capture much of what is valuable about both the abstract minimalist version of *Inferentialism 2* as well as *Interpretation* discussed above.

7. The Fiction view of Models

In this section we discuss a number of recent attempts to analyse scientific modelling by drawing an analogy with literary fiction. We begin by introducing the leading ideas and differentiating between different strands of argument. We then examine a number of accounts that analyse epistemic representation against the backdrop of literary fiction. We finally discuss criticisms of the fiction view.

7.1 Models and Fiction

Scientific discourse is rife with passages that appear to be descriptions of systems in a particular discipline, and the pages of textbooks and journals are filled with discussions of the properties and the behaviour of those systems. Students of mechanics investigate at length the dynamical properties of a system consisting of two or three spinning spheres with homogenous mass distributions gravitationally interacting only with each other. Population biologists study the evolution of one species that reproduces at a constant rate in an unchanging environment. And when studying the exchange of goods, economists consider a situation in which there are only two goods, two perfectly rational agents, no restrictions on available information, no transaction costs, no money, and dealings are done immediately. Their surface structure notwithstanding, no one would mistake descriptions of such systems as descriptions of an *actual* system: we know very well that there are no such systems (of course some models are actual systems – a scale model of a car in a wind tunnel for example – but in this section we focus on models that are not of this kind). Scientists sometimes express this fact by saying that they talk about ‘model-land’ (see, for instance, (Smith 2007, p. 135)).

Thomson-Jones (2010, p. 284) refers to such a description as a ‘description of a missing system’. These descriptions are embedded in what he calls the ‘face value practice’ (ibid. p. 285): the practice of talking and thinking about these systems as if they were real. We observe that the amplitude of an ideal pendulum remains constant over time in much the same way in which we say that the moon’s mass is approximately 7.34×10^{22} kg. Yet the former statement is about a point mass suspended from a massless string – and there is no such thing in the world.

The face value practice raises a number of questions. What account should be given of these descriptions and what sort of objects, if any, do they describe? How should we analyse the face value practice? Are we putting forward truth-evaluable claims when putting forward descriptions of missing systems? An answer to these questions emerges from the following passage by Peter Godfrey-Smith:

‘[...] I take at face value the fact that modelers often *take* themselves to be describing imaginary biological populations, imaginary neural networks, or imaginary economies. [...] Although these imagined entities are puzzling, I suggest that at least much of the time they might be treated as similar to something that we are all familiar with, the imagined objects of literary fiction. Here I have in mind entities like Sherlock Holmes’ London, and Tolkein’s Middle Earth. [...] the model systems of science often work similarly to these familiar fictions.’ (2006, p. 735)

This is the core of the fiction view of models: models are akin to places and characters in literary fiction. When modelling the solar system as consisting of ten perfectly spherical spinning tops physicists describe (and *take themselves* to describe) an imaginary physical system; when considering an ecosystem with only one species biologists describe an imaginary population; and when investigating an economy without money and transaction costs economists describe an imaginary economy. These imaginary scenarios are tellingly like the places and characters in works of fiction like *Madame Bovary* and *Sherlock Holmes*.

Although hardly at the centre of attention, the parallels between certain aspects of science and literary fiction have not gone unnoticed. Maxwell discussed in great detail the motion of a ‘purely imaginary fluid’ in order to understand the electromagnetic field (Niven 1965, pp. 159-160). The parallel between science and fiction occupied centre stage in Vaihinger’s (1911/1924) philosophy of the ‘as if’. More recently, the parallel has also been drawn specifically between models and fiction. Cartwright observes that ‘a model is a work of fiction’ (1983, p. 153) and later suggests an analysis of models as fables (1999, Ch. 2). McCloskey (1990) emphasises the importance of narratives and stories in economics. Fine notes that modelling natural phenomena in every area of science involves fictions in Vaihinger’s sense (Fine 1993, p. 16), and Sklar highlights that describing systems ‘as if’ they were systems of some other kind is a royal route to success (2000, p. 71). Elgin (1996, Ch. 6) argues that

science shares important epistemic practices with artistic fiction. Hartmann (1999) and Morgan (2001) emphasise that stories and narratives play an important role in models, and Morgan (2004) stresses the importance of imagination in model building. Sugden (2000) points out that economic models describe ‘counterfactual worlds’ constructed by the modeller. Frigg (2003) and (2010) suggests that models are imaginary objects, and Grüne-Yanoff and Schweinzer (2008) emphasise the importance of stories in the application of game theory. Toon (2010) and (2012) has formulated an account of representation based on a theory of literary fiction. Contessa (2010) provides a fictional ontology of models and Levy (2012), (2015) discusses models as fictions.

But simply likening modelling to fiction does not solve philosophical problems. Fictional discourse and fictional entities face well-known philosophical questions, and hence explaining models in terms of fictional characters seems to amount to little more than to explain *obscurum per obscurius*. The challenge for proponents of the fiction view is to show that drawing an analogy between models and fiction has heuristic value.

A first step towards making the analogy productive is to get clear on what the problem is that the appeal to fiction is supposed to solve. This issue divides proponents of the fiction view into two groups. Authors belonging to the first camp see the analogy with fiction as providing an answer to the Problem of Ontology. Models, on that view, are *ontologically* on par with literary fiction while there is no productive parallel between models and fiction as far as the ER-Problem (or indeed any other problem of representation) is concerned. Authors belonging to the second group hold the opposite view. They see the analogy with fiction first and foremost as providing an answer to the ER-Problem (although, as we have seen, this may place restrictions on the ontology of models). Scientific representation, on this view, has to be understood along the lines of how literary fiction relates to reality. Positions on ontology vary. Some authors in this group also adopt a fiction view of ontology; some remain agnostic about the analogy’s contribution to the matters of ontology; and some reject the Problem of Ontology altogether.

This being a review of models and representation, we refer the reader to Gelfert's contribution to this book for an in-depth discussion of the ontology of models (Gelfert forthcoming), and focus on the fiction view's contribution to semantics. Let us just note that those who see fiction as providing an ontology of models are spoilt for choice. In principle every option available in the extensive literature on fiction is a candidate for an ontology of models; for reviews of these options see Friend's (Friend 2007) and Salis' (Salis). Different authors have made different choices, with proposals being offered by Contessa (2010), Ducheyne (2008), Frigg (2010), Godfrey-Smith (2009), Levy (2015), and Sugden (2009). Cat (2012), Liu (2012; 2015), Pincock (2012, Ch. 12), Thomson-Jones (2010) and Toon (2012) offer critical discussions of some of these approaches.

Even if these ontological problems were settled in a satisfactory manner, we would not be home and dry yet. Vorms (2011; 2012) argues that what's more important than the entity itself is the format in which the entity is presented. A fiction view that predominantly focuses on understanding the fictional entities themselves (and, once this task is out of the way, their relation to the real-world targets), misses an important aspect, namely how agents draw inferences from models. This, Vorms submits, crucially depends on the format under which they are presented to scientists, and different formats allow scientists to draw different inferences. This ties in with Knuuttila's insistence that we ought to pay more attention to the 'medium of representation' when studying models (2005; 2011).

One last point stands in need of clarification: the meaning of the term 'fiction'. Setting aside subtleties that are irrelevant to the current discussion, the different uses of 'fiction' fall into two groups: fiction as falsity and fiction as imagination (Frigg 2010). Even though not mutually exclusive, the senses should be kept separate. The first use of 'fiction' characterises something as deviating from reality. We brand Peter's account of events a fiction if he does not report truthfully how things have happened. In the second use, 'fiction' refers to a kind of literature, *literary fiction*. Rife prejudice notwithstanding, the defining feature of literary fiction is not falsity. Neither is everything that is said in, say, a novel untrue (novels like *War and Peace* contain correct historical information); nor does every text containing false reports qualify as fiction (a wrong news report or a faulty documentary do not by that token

turn into fiction – they remain what they are, namely wrong factual statements). What makes a text fictional is the attitude that the reader is expected to adopt towards it. When reading a novel we are not meant to take the sentences we read as reports of fact; rather we are supposed to imagine the events described.

It is obvious from what has been said so far that the fiction view of models invokes the second sense of ‘fiction’. Authors in this tradition do not primarily intend to brand models as false; they aim to emphasise that models are presented as something to ponder. This is not to say the first sense of fiction is irrelevant in science. Traditionally fictions in that sense have been used as calculational devices for generating predictions, and recently Bokulich (2009) emphasised the explanatory function of fictions. The first sense of fiction is also at work in philosophy where antirealist positions are described as fictionalism. For instance, someone is a fictionalist about numbers if she thinks that numbers don’t exist (see Kalderon (2005) for a discussion of several fictionalisms of this kind). Scientific anti-realists are fictionalists about many aspects of scientific theories, and hence Fine characterises fictionalism as an ‘antirealist position in the debate over scientific realism’ (Fine 1998) cf. (Fine; 2009), a position echoed in Winsberg (2009) and Suárez (2009). Morrison (2009) and Purves (2013) and offer critical discussions of this approach, which the latter calls fiction as ‘truth conducive falsehood’ (ibid., p. 236); Woods (2014) offers a critical assessment of fictionalism in general. Although there are interesting discussions to be had about the role that this kind of fictions play in the philosophy of science, it is not our interest here.

7.2 Direct Representation

In this subsection and the next we discuss proposals that have used the analogy between models and fiction to elucidate representation.

Most theories of representation we have encountered so far posit that there are model systems and construe epistemic representation as a relation between two entities, the model system and the target system. Toon calls this the *indirect* view of representation (2012, p. 43); Levy, speaking specifically about the fiction view of models, refers to it as the *whole-cloth fiction* view (2012, p. 741). Indeed, Weisberg

views this indirectness as the defining feature of modelling (2007). This view faces the problem of ontology because it has to say what kind of things model systems are. This view contrasts with what Toon (, p. 43) and Levy (2015, p. 790) call a *direct* view of representation (Levy (2012, p. 741) earlier also referred to it as the *worldly fiction* view). This view does not recognise model systems and aims instead to explain epistemic representation as a form of direct description. Model descriptions (like the description of an ideal pendulum) provide an ‘imaginative description of real things’ (Levy 2012, p. 741) such as actual pendula, and there is no such thing as a model system of which the pendulum description is literally true (Toon 2012, pp. 43-44). In what follows we use Toon’s terminology and refer to this approach as ‘direct representation’.

Toon and Levy both reject the indirect approach because of metaphysical worries about fictional entities, and they both argue that the direct view has the considerable advantage that it does not have to deal with the vexed problem of the ontology of model systems and their comparison with real things at all. Levy (2015, p. 790) sees his approach as ‘largely complimentary to Toon’s’. So we first discuss Toon’s approach and then turn to Levy’s.

Toon (2010; 2010; 2012) takes as his point of departure Walton’s (1990) theory of representation in the arts. At the heart of this theory is the notion of a game of make-believe. The simplest examples of these games are children’s plays (ibid., p. 11). In one such play we imagine that stumps are bears and if we spot a stump we imagine that we spot a bear. In Walton’s terminology the stumps are *props*, and the rule that we imagine a bear when we see a stump is a *principle of generation*. Together a prop and a principle of generation prescribe what is to be imagined. If a proposition is so prescribed to be imagined, then the proposition is *fictional* in the relevant game. The term ‘fictional’ has nothing to do with falsity; on the contrary, it indicates that the proposition is ‘true in the game’. The set of propositions actually imagined by someone need not coincide with the set of all fictional propositions in game. It could be the case that there is a stump somewhere that no one has seen and hence no one imagines that it’s a bear. Yet the proposition that the unseen stump is a bear is fictional in the game.

Walton considers a vast variety of different props. In the current context two kinds of props are particularly important. The first are objects like statues. Consider a statue showing Napoleon on horseback (Toon 2012, p. 37). The statue is the prop, and the games of make-believe for it are governed by certain principles of generation that apply to statues of this kind. So when seeing the statue we are mandated to imagine, for instance, that Napoleon has a certain physiognomy and certain facial expressions. We are not mandated to imagine that Napoleon was made of bronze, or that he hasn't moved for more than 100 years.

The second important kind of props are works of literary fiction. In this case the text is the prop, which together with principles of generation appropriate for literary fictions of a certain kind, generates fictional truths by prescribing readers to imagine certain things. For instance, when reading *The War of the Worlds* (ibid., p. 39) we are prescribed to imagine that the dome of St Paul's Cathedral has been attacked by aliens and now has a gaping hole on its western side.

In Walton's theory something is a *representation* if it has the social function of serving as a prop in a game of make believe, and something is an *object of a representation* if the representation prescribes us to imagine something about the object (ibid., p. 35, 39). In the above examples the statue and the written text are the props, and Napoleon and St Paul's Cathedral, respectively, are the objects of the representations.

The crucial move now is to say that models are props in games of make believe. Specifically, material models – such as an architectural model of the Forth Road Bridge – are like the statue of Napoleon (ibid., p. 37): the model is the prop and the Bridge is the object of the representation. The same observation applies to theoretical models, such as a mechanical model of a bob bouncing on a spring. The model portrays the bob as a point mass and the spring as perfectly elastic. The model description represents the real ball and spring system in the same way in which a literary text represents its objects (ibid., pp. 39-40): the model description prescribes imaginings about the real system – we are supposed to imagine the real spring as perfectly elastic and the bob as a point mass.

We now see why Toon's account is a direct view of modelling. Theoretical model descriptions represent actual concrete objects: the Forth Road Bridge and the bob on a spring. There is no intermediary entity of which model descriptions are literally true and which are doing the representing. Models prescribe imaginings about a real world target, and that is what representation consists in.

This is an elegant account of representation, but it is not without problems. The first issue is that it does not offer an answer to the ER-Problem. Imagining that the target has a certain feature does not tell us how the imagined feature relates to the properties the target actually has, and so there is no mechanism to transfer model results to the target. Imagining the pendulum bob to be a point mass tells us nothing about which, if any, claims about point masses are also true of the real bob. Toon mentions this problem briefly. His response is that

‘principles of generation often link properties of models to properties of the system they represent in rather direct way. If the model has a certain property then we are to imagine that system does too. If the model is accurate, then the model and the system will be similar in this respect. [...] [But] not all principles of generation are so straightforward. [...] In some cases similarity seems to play no role at all.’ (ibid., pp. 68-69)

In as far as the transfer mechanism is similarity, the view moves close to the similarity view, which brings with it both some of the benefits and the problems we have discussed in Section 4. The cases in which similarity plays no role are left unresolved and it remains unclear how surrogative reasoning with such models is supposed to happen.

The next issue is that not all models have a target system, which is a serious problem for a view that analyses representation in terms of imagining something *about* a target. Toon is well aware of this issue and calls them *models without objects* (ibid., p. 76). Some of these are models of discredited entities like the ether and phlogiston, which were initially thought to have a target but then turned out not to have one (ibid., p. 76). But not all models without objects are errors: architectural plans of buildings that are never built or models of experiments that are never carried out fall into the same category (ibid., p. 76).

Toon addresses this problem by drawing another analogy with fiction. He points out that not all novels are like *The War of the Worlds*, which has an object. Passages from *Dracula*, for instance, ‘do not represent any actual, concrete object but are instead about fictional characters’ (ibid., p. 54). Models without a target are like passages from *Dracula*. So the solution to the problem is to separate the two cases neatly. When a model has target then it represents that target by prescribing imaginings about the target; if a model has no target it prescribes imaginings about a fictional character (ibid., p. 54).

Toon immediately admits that models without targets ‘give rise to all the usual problems with fictional characters’ (ibid., p. 54). However, he seems to think that this is a problem we can live with because the more important case is the one where models do have a target, and his account offers a neat solution there. He offers the following summative statement of his account (ibid., p. 62):

Direct Representation: A scientific model M represents a target system T iff M functions as prop in game of make-believe.

This definition takes it to be understood that the imaginings prescribed are about the target T if there is a target, and about a fictional character if there isn’t because there need not be any object that the model prescribes imaginings about’ (ibid., p. 81).

This bifurcation of imaginative activities raises questions. The first is whether the bifurcation squares with the face value practice. Toon’s presentation would suggest that the imaginative practices involved in models with targets are very different from the ones involved in models without them. Moreover, they require a different analysis because imagining something about an existing object is different from imagining something about a fictional entity. This, however, does not seem to sit well with scientific practice. In some cases we are mistaken: we think that the target exists but then find out that it doesn’t (as in the case of phlogiston). But does that make a difference to the imaginative engagement with a phlogiston model of combustion? Even today we can understand and use such models in much the same way as its original protagonists did, and knowing that there is no target seems to make little, if

any, difference to our imaginative engagement with the model. Of course the presence or absence of a target matters to many other issues, most notably surrogative reasoning (there is nothing to reason about if there is no target!), but it seems to have little importance for how we imaginatively engage with the scenario presented to us in a model.

In other cases it is simply left open whether there is target when the model is developed. In elementary particle physics, for instance, a scenario is often proposed simply as a suggestion worth considering and only later, when all the details are worked out, the question is asked whether this scenario bears an interesting relation to what happens in nature, and if so what the relation is. So, again, the question of whether there is or isn't a target seems to have little, if any, influence on the imaginative engagement of physicists with scenarios in the research process. This does not preclude different philosophical analyzes being given of modeling with and without a target, but any such analysis will have to make clear the commonalities between the two.

Let us now turn to a few other aspects of *Direct Representation*. The view successfully solves the problem of asymmetry. Even if it uses similarity in response to the ER-Problem, the imaginative process is clearly directed towards the target. An appeal to imagination also solves the problem of misrepresentation because there is no expectation that our imaginations are correct when interpreted as statements about the target. Given its roots in a theory of representation in art, it's natural to renounce any attempts to demarcate scientific representation from other kinds of representation (ibid., p. 62). The Problem of Ontology is dispelled for representations with an object, but it remains unresolved for representations without one. However, *Direct Representation* offers at best a partial answer to the ER-Problem, and nothing is said about either the Problem of Style and or Standards of Accuracy. Similarly, Toon remains silent about the applicability of mathematics.

Levy also rejects an indirect view primarily because of the unwieldiness of its ontology and endorses a direct view of representation (2012, pp. 744-747; 2015, pp. 780-790). Like Toon, he develops his version of the direct view by appeal to Walton's notion of prop oriented make believe. When, for instance, we're asked where in Italy

the town of Crotona lies, we can be told that it's in the arch of the Italian boot. In doing so we are asked to imagine something about the shape of Italy and this imagination is used to convey geographical information. Levy then submits that 'we treat models as games of prop oriented make-believe' (2012, p. 791). Hence modelling consists in imagining something directly about the target.

Levy pays careful attention to the ER-Problem. In his (2012, p. 744) he proposed that the problem be conceptualised in analogy with metaphors, but immediately added that this was only a beginning which requires substantial elaboration. In his (2015, pp. 792-796) he takes a different route and appeals to Yablo's (2014) theory of partial truth. The core idea of this view is that a statement is partially true 'if it is true when evaluated only relative to a subset of the circumstances that make up its subject matter – the subset corresponding to the relevant content-part' (Levy 2015, p. 792). Levy submits that this will also work for a number of cases of modelling, but immediately adds that there are other sorts of cases that don't fit the mould (ibid., p. 794). Such cases often are ones in which distortive idealisations are crucial and cannot be set aside. These require a different treatment and it's an open question what this treatment would be.

Levy offers a radical solution to the problem of models without targets: there aren't any! He first broadens the notion of a target system, allowing for models that are only loosely connected to targets (ibid., pp. 796-797). To this end he appeals to Godfrey-Smith's notion of 'hub-and-spoke' cases: families of models where only some have a target (which makes them the hub models) and the others are connected to them via conceptual links (spokes) but don't have specific target. Levy points out that such cases should be understood as having a *generalised target*. If something that looks like a model doesn't meet the requirement of having even a generalised target, then it's not a model at all. Levy mentions structures like the game of life and observes that they are 'bits of mathematics' rather than models (ibid., p. 797). This eliminates the need for fictional characters in the case of targetless models.

This is a heroic act of liberation, but questions about it remain. The direct view renders fictional entities otiose by positing that a model is nothing but an act of imagining something about a concrete actual thing. But generalised targets are not

concrete actual things, and often not even classes of such things. There is a serious question whether one can still reap the (alleged) benefits of a view that analyses modelling as imaginings about concrete things, if the things about which we imagine something are no longer concrete. Population growth or complex behaviour are not concrete things like rabbits and stumps, and this would seem to pull the rug from underneath a direct approach to representation. Likewise, the claim that models without target are 'just mathematics' stands in need of further elucidation. Looking back at Toon's examples of such models, a view that considers them just mathematics does not come out looking very natural.

7.3 Parables and Fables

Cartwright (2010) focuses on highly idealised models such as Schelling's model of social segregation (1978) and Pissarides' model of the labour market (1992). The problem with these models is that the objects and situations we find in such models are not at all like the things in the world that we are interested in. Cities aren't organised as checkerboards and people don't move according to simple algorithmic rules (as they do in Schelling's model), and there are no labourers who are solely interested in leisure and income (as is the case in Pissarides' model). Yet we are supposed to learn something about the real world from these models. The question is how.

Cartwright submits that an answer to this question emerges from a comparison of models with narratives, in particular fables and parables. An example of a fable is the following: 'A marten eats the grouse; A fox throttles the marten; the tooth of the wolf, the fox. Moral: the weaker are always prey to the stronger.' (ibid., p. 20) The characters in the fable are highly idiosyncratic, and typically we aren't interested in them *per se* - we don't read fables to learn about foxes and martens. What we are interested in is the fable's general and more abstract conclusion, in the above example that the weaker are always prey to the stronger. In the case of the fable the moral is typically built in the story and explicitly stated (ibid.).

Cartwright then invites us to consider the parable of the labourers in the vineyard told in the Gospel of Matthew (ibid.). A man goes to the market to hire day labourers. He

hires the first group early in the morning, and then returns several times during the day to hire more labourers, and he hires the last group shortly before dusk. Some worked all day, while some hardly started when the day ended. Yet he pays the same amount to all of them. Like in a fable, when engaging with a parable the reader takes no intrinsic interest in the actors and instead tries to extract a more general moral. But unlike in fables, in parables no moral appears as part of the parable itself (ibid., p. 29). Hence parables need interpretation, and alternative interpretations are possible. The above fable is often interpreted as being about the entry to God's kingdom, but, as Cartwright observes, it can just as well be interpreted as making the market-based capitalist point that you get what you contract for, and should not appeal to higher forms of justice (ibid., p. 21).

These are features models share with fables and parables: 'like the characters in the fable, the objects in the model are highly special and do not in general resemble the ones we want to learn about' (ibid., p. 20) and the 'lesson of the model is, properly, more abstract than what is seen to happen in the model' (ibid., p. 28). This leaves the question whether models are fables or parables. Some models are like fables in that they have the conclusion explicitly stated in them. But most models are like parables (ibid. p. 29): their lesson is not written in the models themselves (ibid., p. 21), and worse: 'a variety of morals can be attributed to the models' (ibid., p. 21). A model, just like a parable, is interpreted against a rich background of theory and observation, and the conclusion we draw depends to a large extent on the background (ibid., p. 30).

So far the focus was on deriving a conclusion about the model *itself*. Cartwright is clear that one more step is needed: 'in many cases we want to use the results of these models to inform our conclusions about a range of actually occurring (so-called *target*) situations' (ibid., p. 22 original emphasis). In fact, making this transfer of model results to the real world is the ER-Problem. Unfortunately she does not offer much by way of explaining this step and merely observes that 'a description of what happens in the model that does not fit the target gets recast as one that can' (ibid., p. 20). This gestures in the right direction, but more would have to be said about how exactly a model description is recast to allow for transfer of model results to target systems. In earlier work Cartwright observed that what underlies the relationship between models and their targets is a 'loose notion of resemblance' (Cartwright 1999,

pp. 192-193) *cf.* (Cartwright 1999, pp. 261-262). This could be read as suggesting that she would endorse some kind of similarity view of representation. Such a view, however, is independent of an appeal to fables and parables.

In passing we would like to mention that the same kind of models is also discussed in Sugden's (2000; 2009). However, his interest is in induction rather than representation, and if reframed in representational terms then his account becomes a similarity account like Giere's. See Grüne-Yanoff's (2009) and Knuuttila's (2009) for a discussion.

7.4 Against Fiction

The criticisms we have encountered above were intrinsic criticisms of particular versions of the fiction view, and as such they presuppose a constructive engagement with the view's point of departure. Some critics think that any such engagement is misplaced because the view got started on the wrong foot entirely. There are five different lines of attack. The first criticism is driven by philosophical worries about fiction. Fictions, so the argument goes, are intrinsically dubious and are beset with so many serious problems that one should steer away from them whenever possible. So it could be claimed that assigning them a central role in science is a manifestation of philosophical masochism. This, however, overstates the problems with fictions. Sure enough, there is controversy about fictions. But the problems pertaining to fictions aren't more devastating than those surrounding other items on the philosophical curriculum, and these problems surely don't render fictions off limits.

The second criticism, offered for example by Giere (2009, p. 257), is that the fiction view – involuntarily – plays into the hands of irrationalists. Creationists and other science sceptics will find great comfort, if not powerful rhetorical ammunition, in the fact that philosophers of science say that scientists produce fiction. This, so the argument goes, will be seen as a justification of the view that religious dogma is on par with, or even superior to, scientific knowledge. Hence the fiction view of models undermines the authority of science and fosters the cause of those who wish to replace science with religious or other unscientific worldviews.

Needless to say, we share Giere's concerns about creationism. In order not to misidentify the problem it is important to point out that Giere's claim is not that the view itself – or its proponents – support creationism; his worry is that the view is a dangerous tool when it falls into the wrong hands. What follows from this, however, is not that the fiction view itself should be abandoned; but rather that some care is needed when dealing with the press office. As long as the fiction view of models is discussed in informed circles, and, when popularised, is presented carefully and with the necessary qualifications, it is no more dangerous than other ideas, which, when taken out of context, can be put to uses that would (probably) send shivers down the spines of their progenitors (think, for instance, of the use of Darwinism to justify eugenics).

The third objection, also due to Giere, has it that the fiction view misidentifies the aims of models. Giere agrees that from an *ontological* point of view scientific models and works of fictions are on par, but emphasises that '[i]t is their differing function in practice that makes it inappropriate to regard scientific models as works of fiction' (Giere 2009, p. 249). Giere identifies three functional differences (Giere 2009, pp. 249-252). First, while fictions are the product of a single author's individual endeavours, scientific models are the result of a public effort because scientists discuss their creations with their colleagues and subject them to public scrutiny. Second, there is a clear distinction between fiction and non-fiction books, and even when a book classified as non-fiction is found to contain false claims, it is not reclassified as fiction. Third, unlike works of fiction, whose prime purpose is to entertain (although some works can also give insight into certain aspects of human life), scientific models are representations of certain aspects of the world.

These observations, although correct in themselves, have no force against the fiction view of models. First, whether a fiction is the product of an individual or a collective effort has no impact on its status as a fiction; a collectively produced fiction is just a different kind of fiction. Even if *War and Piece* (to take Giere's example) had been written in a collective effort by all established Russian writers of Tolstoy's time, it would still be a fiction. Vice versa, even if Newton had never discussed his model of the solar system with anybody before publishing it, it would still be science. The history of production is immaterial to the fictional status of a work. Second, as we

have seen in subsection 7.1, falsity is not a defining feature of fiction. We agree with Giere that there is a clear distinction between texts of fiction and non-fiction, but we deny that this distinction is defined by truth or falsity; it is the attitude that we are supposed to adopt towards the text's content that makes the difference. Once this is realised, the problem fades away. Third, many proponents of the fiction view (those belonging to the first group mentioned in subsection 7.1) are clear that problems of ontology should be kept separate from function and agree that it is one of the prime function of models to represent. This point has been stressed by Godfrey-Smith (2009, pp. 108-111) and it is explicit in other views such as Frigg's (2010).

The fourth objection is due to Magnani, who dismisses the fiction view for misconstruing the role of models in the process of scientific discovery. The fundamental role played by models, he emphasises, 'is the one we find in the core conceptual discovery processes, and that these kinds of models cannot be indicated as fictional at all, because they are constitutive of new scientific frameworks and new empirical domains.' (2012, p. 3) This criticism seems to be based on an understanding of fiction as falsity because falsities can't play a constitutive role in the constitution of new empirical domains. We re-iterate that the fiction view is not committed to the 'fiction as falsity' account and hence is not open to this objection.

The fifth objection is that fictions are superfluous and hence should not be regarded as forming part of (let alone *being*) scientific models because we can give a systematic account of how scientific models work without invoking fictions. This point has been made in different ways by Pincock (2012, Ch. 12) and Weisberg (2013, Ch. 4) (for a discussion of Weisberg's arguments see Odenbaugh (2015)). We cannot do justice to the details of their sophisticated arguments here, and will concern ourselves only with their main conclusion. They argue that scientific models are mathematical objects and that they relate to the world due to the fact that there is a relationship between the mathematical properties of the model and the properties found in the target system (in Weisberg's version similarity relations to a parametrised version of the target). In other words, models are mathematical structures and they represent due to there being certain mathematical relations between these structures and a mathematical rendering of the target system. (Weisberg includes fictions as convenient 'folk ontology' that may serve as a crutch when thinking about the model, but takes them to be ultimately

dispensable when it comes to explaining how models relate to the world). This, however, brings us back to a structuralist theory of representation, and this theory, as we have seen in Section 5, is far from unproblematic. So it is at best an open question whether getting rid of fiction provides an obvious advantage.

8. Representation-As

In this section we discuss approaches that depart from Goodman's notion of 'representation-as' (Goodman 1976). In his account of aesthetic representation the idea is that a work of art does not just denote its subject, but moreover it represents it as being thus or so. Elgin (2010) further developed this account and, crucially, suggested that it also applies to scientific representations. This is a vital insight and it provides the entry point to what we think of as the most promising account of epistemic representation.

In this section we present Goodman and Elgin's notion of 'representation as', and outline how it is a complex type of reference involving a mixture of denotation and what they call 'exemplification'. We introduce the term of art 'representation-as' to indicate that we are talking about the specific concept that emerges from Goodman's and Elgin's writings. We then discuss how the account needs to be developed in the context of scientific representation. And finally we present our own answer to the ER-Problem, and demonstrate how it answers the questions laid out in Section 2.

8.1 Exemplification and Representation-As

Many instances of epistemic representation are instances of representation-as. Caricatures are paradigmatic examples: Churchill is represented as a bulldog, Thatcher is represented as a boxer, and the Olympic Stadium is represented as a UFO. Using these caricatures we can attempt to learn about their targets: attempt to learn about a politician's personality or a building's appearance. The notion applies beyond caricatures. Holbein's *Portrait of Henry VIII* represents Henry as imposing and powerful and Stoddart's statue of David Hume represents him as thoughtful and wise. The leading idea is that scientific representation works in much the same way. A

model of the solar system represents the sun as perfect sphere; the logistic model of growth represents the population as reproducing at fixed intervals of time; and so on. In each instance, models can be used to attempt to learn about their targets by determining what the former represent the latter as being. So representation-as relates, in a way to be made more specific below, to the Surrogative Reasoning Condition discussed in Section 2.

The locution of representation-as functions in the following way: an object X (e.g. a picture, statue, or model) represents a subject Y (e.g. a person or target system) as being thus or so (Z). The question then is what establishes this sort of representational relationship. The answer requires presenting some of the tools Goodman and Elgin use to develop their account of representation-as.

One of the central posits of Goodman's account is that denotation is 'the core of representation' (Goodman 1976, p. 5). Stoddart's statue of David Hume denotes Hume and a model of the solar system denotes the solar system. In that sense the statue and the model are representations *of* their respective targets. To distinguish representation of something from other notions of representation we introduce the technical term 'representation-of'. Denotation is what establishes representation-of. (For a number of qualifications and caveats about denotation see our (forthcoming, Sec. 2)).

Not all representations are a representation-of. A picture showing a unicorn is not a representation-of a unicorn because things that don't exist can't be denoted. Yet there is a clear sense in which such a picture is a representation. Goodman and Elgin's solution to this is to distinguish between being a representation-of something and being a something-representation (Elgin 2010, pp. 1-2; Goodman 1976, pp. 21-26). What makes a picture a something-representation (despite the fact it may fail to denote anything) is that it is the sort of symbol that denotes. Elgin argues:

'A picture that portrays a griffin, a map that maps the route to Mordor [...] are all representations, although they do not represent anything. *To be a representation, a symbol need not itself denote, but it needs to be the sort of symbol that denotes.* Griffin pictures are representations then because they are animal pictures, and some

animal pictures denote animals. Middle Earth maps are representations because they are maps and some maps denote real locations. [...] *So whether a symbol is a representation is a question of what kind of symbol it is.*' (Elgin 2010, pp. 1-2, emphasis added)

These representations can be classified into genres, in a way that does not depend on what they are representations-of (since some may fail to denote), but instead on what they portray. In the case of pictures, this is fairly intuitive (how this is to be developed in the case of scientific models is discussed below). If a picture portrays a man, it is a man-representation, if it portrays a griffin it is a griffin-representation and so on. In general, a picture *X* is *Z*-representation if it portrays *Z*. The crucial point is that this does not presuppose that *X* be a representation-of *Z*; indeed *X* can be *Z*-representation without denoting anything. A picture must denote a man to be a representation-of a man. But it need not denote anything to be a man-representation.

The next notion we need to introduce is *exemplification*. An item exemplifies a property if it at once instantiates the property and refers to it:

‘Exemplification is possession plus reference. To have without symbolising is merely to possess, while to symbolise without having is to refer in some other way than by exemplifying’ (Goodman 1976, p. 53).

Exemplification is a mode of reference that holds between items and properties. In the current context properties are to be understood in the widest possible sense. An item can exemplify one-place properties, multi-place properties (i.e. relations), higher order properties, structural properties, etc. Paradigmatic examples of exemplification are samples. A chip of paint on a manufacturer’s sample card both instantiates a certain colour, and at the same time refers to that colour (Elgin 1983, p. 71).

But although exemplification requires instantiation, not every property instantiated by an object is exemplified by it. The chip of paint does not, for example, exemplify its shape or its location on the card. In order to exemplify a property, an object must both instantiate the property and the property itself must be made epistemically salient.

How saliency is established will be determined on a case-by-case basis, and we say more about this below.

We can now turn to the conditions under which X represents Y as Z . A first stab would be to say that X represents Y as Z if X is a Z -representation and denotes Y . This however, is not yet good enough. It is important that properties of Z are ‘transferred’ to Y . Elgin makes this explicit:

‘[X] does not merely denote [Y] and happen to be a [Z]-representation. Rather in being a [Z]-representation, [X] exemplifies certain properties and *imputes* those properties or related ones to [Y]. [...] The properties exemplified in the [Z]-representation thus serve as a bridge that connects [X] to [Y].’(2010, p. 10, emphasis added)

This gives a name to the crucial step: imputation. This step can be analysed in terms of stipulation by a user of a representation. When someone uses X as a representation-as, she has to stipulate that certain properties that are exemplified in X be imputed to Y . We emphasise that imputation does not imply truth: Y may or may not have the properties imputed to it by X . So the representation can be seen as generating a claim about Y that can be true or false; it should not be understood as producing truisms.

Applied to scientific models, the account of epistemic representation that emerges from Goodman and Elgin’s discussion of representation can then be summarised as follows:

Representation-As: A scientific model M represents a target system T iff (i) M denotes T , (ii) M is a Z -representation exemplifying properties P_1, \dots, P_n , and (iii) P_1, \dots, P_n , or related properties, are imputed to T .

It should be added that the first condition can easily be extended to include part-part denotation. In a family portrait the entire portrait denotes the family; at the same time a part of the portrait can denote the mother and another part the father. This is obvious and unproblematic.

We think that this account is on the right track, but all three conditions need to be further developed to furnish a full-fledged account of epistemic representation (at least as applied to scientific models. The developments needed are of different kinds, though. The first condition needs more specificity. How is denotation characterised? What different ways of establishing denotation are there? And how is denotation established in particular cases? These are but some of the questions that a complete account of epistemic representation will have answer. In many cases epistemic representation seems to borrow denotation from linguistic descriptions in which they are embellished and denotation is in effect borrowed from language. So the philosopher of science can turn to the philosophy of language to get a deeper understanding of denotation. This is an interesting project, but it is not one we can pursue here.

In contrast with denotation the other two conditions need to be reformulated because an account moulded on visual representations is only an imperfect match for scientific representations. This is the task for the next section.

8.2 From Pictures to Models: The DEKI Account

According to Goodman and Elgin, for a picture to be a *Z*-representation it has to be the kind of symbol that denotes. On the face of it, there is a mismatch between pictures and scientific models in this regard. The Schelling model represents social segregation with a checkerboard; billiard balls are used to represent molecules; the Phillips-Newlyn model uses a system of pipes and reservoirs to represent the flow of money through an economy; and the worm *Caenorhabditis elegans* is used as a model of other organisms. But neither checkerboards, billiard balls, pipes, or worms seem to belong to classes of objects that typically denote. The same observation applies to scientific fictions (frictionless planes, utility maximising agents, and so on) and the mathematical objects used in science. In fact, matrices, curvilinear geometries, Hilbert spaces etc. were all studied as mathematical objects before they became important in the empirical sciences.

Rather than relying on the idea that scientific models belong to classes of objects that typically denote we propose directly introducing an agent and ground representation

in this agent's actions. Specific checkerboards, systems of pipes, frictionless places and mathematical structures, are epistemic representations because they are used by an agent to represent a system. When an agent uses an object as a representation, we call it a 'base'.

What allows us to classify bases into *Z*-representations is also less clear in the case of scientific representation. We approach this issue in two steps. The first is to recognise the importance of the intrinsic constitution of the base. Pictures are typically canvasses covered with paint. They are classified as *Z*-representations because under appropriate circumstances the canvass is recognised as portraying a *Z*. Much can be said about the canvass' material constitution (the thickness or chemical constitution of the paint, etc.), but these are generally of little interest to understanding what the picture portrays. By contrast, the properties of a scientific model – *qua* material object – do matter. How water flows through the pipes in the Phillips-Newlyn model is crucial to how it represents the movement of money in an economy. That the *Caenorhabditis elegans* is a biological organism is of vital importance for how it is used representationally. In fact, models are frequently classified according to what their material base is. We talk about a pipe model of the economy or worm model of cell division because their bases are pipes and worms. Here we introduce a term of art to recognise that scientific models are generally categorised according to their material constitution. An *O*-object specifies the kind of object something is, *qua* physical object.

O-objects become representations when they are used as such. But how are they classified as *Z*-representations? How does the Phillips-Newlyn machine become an economy-representation, or how does a collection of billiard balls become a gas-representation? (Again, recall that this is not because they denote economies or gases). We suggest, and this is the second step, that this requires an act of *interpretation* (notice that we do not use 'interpretation' in the same sense as Contessa). In the case of pictures, the nature of this interpretation has been the centre of attention for a good while: how one sees a canvass covered with paint as showing a cathedral is regarded by many as one of the important problems of aesthetics. Schier (1986, p. 1) dubbed it the 'enigma of depiction', and an entire body of literature is concerned with it (Kulvicki (2006) provides a useful review). In the case of

scientific models we don't think a simple and universal account of how models are interpreted as *Z*-representations can be given. Interpreting an *O*-object as a *Z*-representation requires attributing properties of *Z*s to the object. How this is done will depend on disciplinary traditions, research interests, background theory and much more. In fact, 'interpretation' is a blank to be filled, and it will be filled differently in different cases.

Some examples should help elucidate what we mean by this. In the case of scale models the interpretation is 'close' to the *O*-object in that it interprets the object in its 'own' terms. The small car is interpreted as a car-representation and the small ship is interpreted as a ship-representation. Likewise, in the case of the Army Corps' model of the San Francisco Bay (Weisberg 2013), parts of the model bay are interpreted in terms of the real bay. In cases like these, the same predicates that apply to the base (qua *O*-object) are applied to the object in order to make it into a *Z*-representation (here $O=Z$). But this is not always the case. For example, the Phillips-Newlyn machine is a system of pipes and reservoirs, but it becomes an economy-representation only when the quantity and flow of water throughout the system are interpreted as the quantity and flow of money throughout an economy. The system is interpreted in terms of predicates that do not apply to the object (qua *O*-object), but turn it into a *Z*-representation (here *O* and *Z* come apart). In sum, an *O*-object that has been chosen as the base of a representation becomes a *Z*-representation if *O* is interpreted in terms of *Z*.

Next in line is exemplification. Much can be said about exemplification in general, but the points by and large carry over from the general discussion to the case of models without much ado. There is one difference, though, in cases like the Phillips-Newlyn machine. Recall that exemplification was defined as the instantiation of a property *P* by an object in such a way that the object thereby refers to *P*. How can the Phillips-Newlyn machine exemplify economic properties when it does not, strictly speaking, instantiate them? The crucial point is that nothing in the current account depends on instantiation being literal instantiation. On this point we are in agreement with Goodman and Elgin, whose account relies on non-literal instantiation. The portrait of Henry cannot, strictly speaking, instantiate the property of being male,

even if it represents him as such. Goodman and Elgin call this metaphorical instantiation (Elgin 1983, p. 81; Goodman 1976, pp. 50-51).

What matters is that properties are epistemically accessible and salient, and this can be achieved with what we call *instantiation-under-an-interpretation I*, *I-instantiation* for short. An economic interpretation of the Phillips-Newlyn machine interprets amounts of water as amounts of money. It does so by introducing a clearly circumscribed rule of proportionality: x litres of water correspond to y millions of the model-economy's currency. This rule is applied without exception when the machine is interpreted as an economy-representation. So we say that under the economic interpretation I_e the machine ' I_e -instantiates' money properties. With the notion of *I-instantiation* at hand, exemplification poses no problem.

The final issue to clear is the imputation of the model's exemplified properties to the target system. In particular, which properties are so imputed? Elgin describes this as the imputation of the properties exemplified by M 'or related ones'. The observation that the properties exemplified by a scientific model and the properties imputed to its target system need not be identical is correct. In fact, few, if any, models in science portray their targets as exhibiting exactly the same features as the model itself. The problem with invoking 'related' properties is not its correctness, but its lack of specificity. Any property can be related to any other property in some way or other, and as long as no specific relation is specified it remains unclear which properties are imputed onto the system.

In the context of science, the relation between the properties exemplified and the ones ascribed to the system is sometimes described as one of simplification (Elgin 1996, p. 184), idealisation (Elgin 1996, p. 184) and approximation (Elgin 2010, p. 11). This could suggest that 'related ones' means 'idealised', at least in the context of science (we are not attributing this claim to Elgin; we are merely considering the option), perhaps similar to the way in which Ducheyne's account discussed above took target properties to be approximations of model properties. But shifting from 'related' to 'idealised' or 'approximated' (or any of their cognates) makes things worse, not better. For one, 'idealisation' can mean very different things in different contexts and hence describing the relation between two properties as 'idealisation' adds little

specificity (see Jones' (2005) for a discussion of different kinds of idealisation). For another, while the relationship between some representation-target properties may be characterised in terms of idealisation, many cannot. A map of the world exemplifies a distance of 29cm between the two points labelled 'Paris' and 'New York'; the distance between the two cities is 5800km; but 29cm is not an idealisation of 5800km. A scale model of a ship being towed through water is not an idealization of an actual ship, at least not in any obvious way. Or in standard representations of Mandelbrot sets the colour of a point indicates the speed of divergence of an iterative function for certain parameter value associated with that point, but colour is not an idealisation of divergence speed.

For this reason it is preferable, in our view, to build a specification of the relationship between model properties and target properties directly into an account of epistemic representation. Let P_1, \dots, P_n be the properties exemplified by M , and let Q_1, \dots, Q_m be the 'related' properties that M imputes to Y (where n and m are positive natural numbers which can but need not be equal). Then the representation M must come with a key K that specifies how exactly P_1, \dots, P_n are converted into Q_1, \dots, Q_m (Frigg 2010). Borrowing notation from algebra (somewhat loosely) we can write $K(\{P_1, \dots, P_n\}) = \{Q_1, \dots, Q_m\}$. K can, but need not be, the identity function; any rule that associates a unique set Q_1, \dots, Q_m with P_1, \dots, P_n is admissible. The relevant clause in the definition of representation-as then becomes: M exemplifies P_1, \dots, P_n and the representation imputes properties Q_1, \dots, Q_m to T where the two sets of properties are connected to each other by a key K .

The above examples help illustrate what we have in mind. Let us begin with the example of the map. (In fact the idea of a key is motivated by a study of maps; for a discussion of maps see Galton's (2001) and Sismondo and Chrisman's (2001).) P is a measured distance on the map between the point labelled 'New York' and the point labelled 'Paris'; Q is the distance between New York and Paris in the world; and K is the scale of the map (in the above case, 1:20,000,000). So the key allows us to translate a property of the map (the 29cm distance) into a property of the world (that New York and Paris are 5800km apart). But the key involved in the scale model of the ship is more complicated. One of the P s in this instance is the resistance the model ship faces when moved through the water in a tank. But this doesn't translate into the

resistance faced by the actual ship in the same way in which distances in a map translate into distances in reality. In fact, the relation between the resistance of the model and the resistance of the real ship stand in a complicated non-linear relationship because smaller models encounter disproportionate effects due to the viscosity of the fluid. The exact form of the key is often highly non-trivial and emerges as the result of a thoroughgoing study of the situation; see Sterrett's (2006) for a discussion of fluid mechanics. In the representation of the Madelbrod set in (Argyris et al. 1994, p. 660) a key is used that translates colour into divergence speed (ibid., p. 695). The square shown is a segment of the complex plane and each point represents a complex number. This number is used as parameter value for an iterative function. If the function converges for number c , then the point in the plane representing c is coloured black. If the function diverges, then a shading from yellow over green to blue is used to indicate the speed of divergence, where yellow is slow, green is in the middle and blue is fast.

Neither of these keys is obvious or trivial. Determining how to move from properties exemplified by models to properties of their target systems can be a significant task, and should not go unrecognized in an account of scientific representation. In general K is a blank to be filled, and it depends on a number of factors: the scientific discipline, the context, the aims and purposes for which M is used, the theoretical backdrop against which M operates, etc. Building K into the definition of representation-as does not prejudge the nature of K , much less single out a particular key as the correct one. The requirement merely is that there must be *some* key for M to qualify as a representation-as.

With these modifications in place we can now formulate our own account of representation (Frigg and Nguyen (forthcoming)). Consider an agent who chooses an O -object as the base of representation and turns it into Z -representation by adopting an interpretation I . Let M refer to the package of the O -object together with the interpretation I that turns it into a Z -representation. Then:

DEKI: A scientific model M represents a target T iff (i) M denotes T (and, possibly, parts of M denote parts of T); (ii) M is a Z -representation exemplifying properties P_1, \dots, P_n ; (iii) M comes with a key, K , specifying how P_1, \dots, P_n are translated into a set

of features $Q_1, \dots, Q_m: K(\{P_1, \dots, P_n\}) = \{Q_1, \dots, Q_m\}$; and (iv) The model imputes at least one of the properties Q_1, \dots, Q_m onto T .

We call this the DEKI account of representation to highlight its key features: denotation, exemplification, keying-up and imputation.

Before highlighting some issues with this account, let us clarify how the account answers the questions we laid out in Section 2. Firstly, as an answer to the ER-Problem, DEKI provides an abstract framework in which to think about epistemic representation. In general, what concretises each of the conditions needs to be investigated on a case-by-case basis. But far from being a defect, this degree of abstractness is an advantage. ‘Epistemic representation’, and even the narrower ‘model-representation’, are umbrella terms covering a vast array of different activities in different fields, and a view that sees representations in fields as diverse as elementary particle physics, evolutionary biology, hydrology and rational choice theory work in exactly the same way is either mistaken or too coarse to make important features visible. DEKI occupies the right middle ground: it is general enough to cover a large array of cases and yet it highlights what all instances of scientific representation have in common. At the same time the account offers an elegant solution to the problem of models without targets: a model that apparently represents Z while there is no Z is a Z -representation but not representation of a Z .

It should be clear how we can use models to perform surrogate reasoning about their targets according to DEKI. The account requires that we investigate the properties that are exhibited by the model. These are then translated into a set of properties that are imputed onto the target. This act of imputation supplies a hypothesis about the target system: does it, or does it not, have those properties? This hypothesis does not have to be true, and as such DEKI allows for the possibility of misrepresentation in a straightforward manner.

DEKI’s abstract character also allows us to talk about different styles of representation. Style, on the DEKI account, is not a monolithic concept; instead it has several dimensions. Firstly, different O -objects can be chosen. In this way we may speak, say, of the *checkerboard style* and of the *cellular automaton style*. In each case

a specific kind of object has been chosen for various modelling purposes. Secondly, the notion of an interpretation allows us to talk about how closely connected the properties of the model are to those that the object *I*-instantiates. Thirdly, different types of keys could be used to characterise different styles. In some instances the key might be the identity key, which would amount to a style of modelling that aims to construct replicas of target systems; in other cases the key incorporates different kinds of idealisations or abstractions, which gives rise to idealisation and abstraction keys. But different keys may be associated with entirely different representational styles.

Similarly, DEKI suggests that there is no significant difference between scientific representations and other kinds of epistemic representation, at least at the general level. However, this is not to say that the two cannot be demarcated whatsoever. The sorts of interpretations under which pictures portray *Z*s seem to be different to the sorts of interpretations that are adopted in the scientific framework. Whether or not this can be cashed out more specifically is an interesting question that we cannot investigate here.

Many details in DEKI still need to be spelled out. But the most significant difficulty, perhaps, arises in connection with the Problem of Ontology. It is not by accident that we have illustrated the account with a physical model, the Phillips-Newlyn machine. Exemplification requires instantiation, which is easily understood for material models, but is highly problematic in the context of nonconcrete models. One option is to view models as fictional entities as discussed in Section 7. But whether, and if so how, fictional entities instantiate properties is controversially discussed and more philosophical work is needed to make sense of such a notion. It is therefore an open question how this account works for nonconcrete models; for a discussion and a proposal see Frigg and Nguyen (2016).

Finally, the account provides us with resources with which to think about the applicability of mathematics. Like the problem of style, various options are available. Firstly, mathematical structures themselves can be taken to be *O*-objects and feature as bases of representation. They can be interpreted on their own terms and therefore exemplify strictly mathematical properties. If one were of a structuralist bent, then the appropriate mathematical properties could be ‘structural’ which could then be

imputed onto the target system (although notice that this approach faces a similar problem to the question of target-end structure discussed in subsection 5.4). Alternatively, the key could provide a translation of these mathematical properties into ones more readily applicable to physical systems. A third alternative would be to take scientific models to be fictional objects, and then adopt an interpretation towards them under which they exemplify mathematical properties. Again, these could be imputed directly onto the target system, or translated into an alternative set of properties. Finally, these fictional models could themselves exemplify physical properties, but in doing so exemplify structural ones as well. Whenever a physical property is exemplified, this provides an extensional relation defined over the objects that instantiate it. The pros and cons of each of these approaches demands further research, but for the purposes of this chapter we simply note that DEKI puts all of these options on the table. Using the framework of *O*-objects, interpretations, exemplification, keys, and imputation provides a novel way in which to think about the applicability of mathematics.

9. Envoi

We reviewed theories of epistemic representation. That each approach faces a number of challenges and that there is no consensus on the matter will not have come as a surprise to anybody. We hope, however, that we managed to map the lay of the land and to uncover the fault lines, and thereby aid future discussions.

Acknowledgements

The authors are listed alphabetically; the chapter is fully collaborative. We would like to thank Demetris Portides and Fiora Salis for helpful comments on an earlier draft.

References

- Abell, Catherine. 2009. "Canny Resemblance." *Philosophical Review* 118:183-223.
- Adams, Ernest W. 1959. "The Foundations of Rigid Body Mechanics and the Derivation of Its Laws from Those of Particle Mechanics " In *The Axiomatic Method: With Special Reference to Geometry and Physics*, ed. Leon Henkin, Patrick Suppes and Alfred Tarski, 250-65. Amsterdam: North-Holland.
- Ainsworth, Peter. 2009. "Newman's Objection." *British Journal for the Philosophy of Science* 60:135-71.
- Ankeny, Rachel A. and Sabina Leonelli. 2011. "What's So Special About Model Organisms?" *Studies in History and Philosophy of Science* 42:313-23.
- Apostel, Leo. 1961. "Towards the Formal Study of Models in the Non-Formal Sciences." In *The Concept and the Role of the Model in Mathematics and Natural and Social Sciences*, ed. Hans Freudenthal, 1-37. Dordrecht: Reidel.
- Argyris, John. H., Gunter Faust and Maria Haase. 1994. *Die Erforschung Des Chaos: Eine Einführung Für Naturwissenschaftler Und Ingenieure*. Braunschweig: Vieweg+Teubner Verlag.
- Armstrong, David M. 1989. *Universals: An Opinionated Introduction* London: Westview Press.
- Aronson, Jerrold L., Rom Harré and Eileen Cornell Way. 1995. *Realism Rescued: How Scientific Progress Is Possible* Chicago: Open Court.
- Bailer-Jones, Daniela M. 2003. "When Scientific Models Represent." *International Studies in the Philosophy of Science* 17:59-74.
- Balzer, Wolfgang, Carles Ulises Moulines and Joseph D. Sneed. 1987. *An Architectonic for Science the Structuralist Program*. Dordrecht: D. Reidel.
- Bartels, Andreas. 2006. "Defending the Structural Concept of Representation." *Theoria* 21:7-19.
- Black, Max. 1973. "How Do Pictures Represent?" In *Art, Perception, and Reality*, ed. Ernst Gombrich, Julian Hochberg and Max Black, 95-130. Baltimore and London: Johns Hopkins University Press
- Bogen, James and James Woodward. 1988. "Saving the Phenomena." *Philosophical Review* 97:303-52.

- Bokulich, Alisa. 2009. "Explanatory Fictions." In *Fictions in Science. Philosophical Essays on Modelling and Idealization*, ed. Mauricio Suárez, 91-109. London and New York: Routledge.
- Bolinska, Agnes. 2013. "Epistemic Representation, Informativeness and the Aim of Faithful Representation." *Synthese* 190:219-34.
- Boniolo, Giovanni. 2007. *On Scientific Representations: From Kant to a New Philosophy of Science*. Hampshire and New York: Palgrave Macmillan.
- Boolos, George S. and Richard C. Jeffrey. 1989. *Computability and Logic*. Cambridge: Cambridge University Press.
- Brading, Katherine and Elaine Landry. 2006. "Scientific Structuralism: Presentation and Representation." *Philosophy of Science* 73:571-81.
- Brandom, Robert B. 1994. *Making It Explicit: Reasoning, Representing and Discursive Commitment*. Cambridge MA: Harvard University Press.
- Brandom, Robert B. 2000. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge MA: Harvard University Press.
- Budd, Malcolm. 1993. "How Pictures Look." In *Virtue and Taste*, ed. D. Knowles and J. Skorupski, 154-75. Oxford: Blackwell.
- Bueno, Otávio. 1997. "Empirical Adequacy: A Partial Structure Approach." *Studies in the History and Philosophy of Science* 28:585-610.
- Bueno, Otávio. 1999. "What Is Structural Empiricism? Scientific Change in an Empiricist Setting." *Erkenntnis* 50:59-85.
- Bueno, Otávio. 2010. "Models and Scientific Representations." In *New Waves in Philosophy of Science*, ed. P.D. Magnus and Jacob Busch, 94-111. Hampshire: Palgrave MacMillan.
- Bueno, Otavio and Mark Colyvan. 2011. "An Inferential Conception of the Application of Mathematics." *Nous* 45:345-74.
- Bueno, Otávio and Steven French. 2011. "How Theories Represent." *British Journal for the Philosophy of Science* 62:857-94.
- Bueno, Otávio, Steven French and James Ladyman. 2002. "On Representing the Relationship between the Mathematical and the Empirical." *Philosophy of Science* 69:452-73.
- Byerly, Henry. 1969. "Model-Structures and Model-Objects." *The British Journal for the Philosophy of Science* 20:135-44.

- Callender, Craig and Jonathan Cohen. 2006. "There Is No Special Problem About Scientific Representation." *Theoria* 55:7-25.
- Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Cartwright, Nancy. 1999. *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Cartwright, Nancy. 1999. "Models and the Limits of Theory: Quantum Hamiltonians and the Bcs Models of Superconductivity." In *Models as Mediators: Perspectives on Natural and Social Science*, ed. Mary Morgan and Margaret Morrison, 241–81. Cambridge: Cambridge University Press.
- Cartwright, Nancy. 2010. "Models: Parables V Fables." In *Beyond Mimesis and Convention. Representation in Art and Science*, ed. Roman Frigg and Matthew C. Hunter, 19-32. Berlin and New York: Springer.
- Cat, Jordi. 2012. "Book Review: Mauricio Suárez (Ed.): *Fictions in Science. Philosophical Essays on Modeling and Idealization*." *Journal for General Philosophy of Science* 43:187-94.
- Chakravartty, Anjan. 2001. "The Semantic or Model-Theoretic View of Theories and Scientific Realism." *Synthese* 127:325-45.
- Chakravartty, Anjan. 2010. "Informational Versus Functional Theories of Scientific Representation." *Synthese* 172:197-213.
- Contessa, Gabriele. 2007. "Scientific Representation, Interpretation, and Surrogate Reasoning." *Philosophy of Science* 74:48–68.
- Contessa, Gabriele. 2010. "Scientific Models and Fictional Objects." *Synthese* 172:215-29.
- Contessa, Gabriele. 2011. "Scientific Models and Representation." In *The Continuum Companion to the Philosophy of Science*, ed. Steven French and Juha Saatsi, 120–37. London: Continuum Press.
- Da Costa, Newton C. A. and Steven French. 1990. "The Model-Theoretic Approach to the Philosophy of Science." *Philosophy of Science* 57:248-65.
- Da Costa, Newton C. A. and Steven French. 2000. "Models, Theories, and Structures: Thirty Years On." *Philosophy of Science* 67:116-27.
- Da Costa, Newton C. A. and Steven French. 2003. *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. Oxford: Oxford University Press.

- de Donato Rodriguez, Xavier and Jesus Zamora Bonilla. 2009. "Credibility, Idealisation, and Model Building: An Inferential Approach." *Erkenntnis* 70:101-18.
- Decock, Lieven and Igor Douven. 2011. "Similarity after Goodman." *Review of Philosophy and Psychology* 2:61-75.
- Demopoulos, William. 2003. "On the Rational Reconstruction of Our Theoretical Knowledge." *British Journal for the Philosophy of Science* 54:371-403.
- Donnellan, Keith S. 1968. "Putting Humpty Dumpty Together Again " *Philosophical Review* 77:203-15.
- Downes, Stephen M. 2009. "Models, Pictures, and Unified Accounts of Representation: Lessons from Aesthetics for Philosophy of Science." *Perspectives on Science* 17:417-28.
- Ducheyne, Steffen. 2008. "Towards an Ontology of Scientific Models." *Metaphysica* 9:119-27.
- Ducheyne, Steffen. 2012. "Scientific Representations as Limiting Cases." *Erkenntnis* 76:73-89.
- Dummett, Michael. 1991. *Frege: Philosophy of Mathematics*. London: Duckworth.
- Elgin, Catherine Z. 1983. *With Reference to Reference*. Indianapolis Hackett.
- Elgin, Catherine Z. 1996. *Considered Judgment*. Princeton: Princeton University Press.
- Elgin, Catherine Z. 2010. "Telling Instances." In *Beyond Mimesis and Convention: Representation in Art and Science*, ed. Roman Frigg and Matthew C. Hunter, 1-18. Berlin and New York: Springer
- Elkins, James. 1999. *The Domain of Images*. Ithaca and London: Cornell University Press.
- Enderton, Herbert B. 1972/2001. *A Mathematical Introduction to Logic*. San Diego and New York: Harcourt.
- Fine, Arthur. 1993. "Fictionalism." *Midwest Studies in Philosophy* 18:1-18.
- Fine, Arthur. 1998. "Fictionalism." In *Routledge Encyclopedia of Philosophy*, ed. Edward Craig, London: Routledge.
- Fine, Arthur. 2009. "Science Fictions: Comment on Godfrey-Smith." *Philosophical Studies* 143:117-25.

- French, Steven. 2000. "The Reasonable Effectiveness of Mathematics: Partial Structures and the Application of Group Theory to Physics." *Synthese* 125:103-20.
- French, Steven. 2003. "A Model-Theoretic Account of Representation (or, I Don't Know Much About Art...But I Know It Involves Isomorphism)." *Philosophy of Science* 70:1472-83.
- French, Steven. 2010. "Keeping Quiet on the Ontology of Models." *Synthese* 172:231-49.
- French, Steven. 2014. *The Structure of the World. Metaphysics and Representation.* Oxford: Oxford University Press.
- French, Steven and James Ladyman. 1999. "Reinflating the Semantic Approach." *International Studies in the Philosophy of Science* 13:103-21.
- French, Steven and Juha Saatsi. 2006. "Realism About Structure: The Semantic View and Nonlinguistic Representations." *Philosophy of Science* 73:548-59.
- French, Steven and Peter Vickers. 2011. "Are There No Things That Are Scientific Theories?" *British Journal for the Philosophy of Science* 62:771-804.
- Friend, Stacie. 2007. "Fictional Characters " *Philosophy Compass* 2:141-56.
- Frigg, Roman. 2002. "Models and Representation: Why Structures Are Not Enough." *Measurement in Physics and Economics Project Discussion Paper Series, DP MEAS 25/02*
- Frigg, Roman. 2003. *Re-Presenting Scientific Representation.* London: PhD: London School of Economics and Political Science
- Frigg, Roman. 2006. "Scientific Representation and the Semantic View of Theories." *Theoria* 55:49-65.
- Frigg, Roman. 2010. "Fiction and Scientific Representation." In *Beyond Mimesis and Convention: Representation in Art and Science*, ed. Roman Frigg and Matthew Hunter, 97-138. Berlin and New York: Springer.
- Frigg, Roman. 2010. "Fiction in Science." In *Fictions and Models: New Essays*, ed. John Woods, 247-87. Munich: Philosophia Verlag.
- Frigg, Roman. 2010. "Models and Fiction." *Synthese* 172 251-68.
- Frigg, Roman, Seamus Bradley, Hailiang Du and Leonard A. Smith. 2014. "Laplace's Demon and the Adventures of His Apprentices." *Philosophy of Science* 81:31-59.

- Frigg, Roman and James Nguyen. 2016. "The Fiction View of Models Reloaded."
Forthcoming in *The Monist*, July 2016.
- Frigg, Roman and James Nguyen. Forthcoming. "Scientific Representation Is
Representation As." In *Philosophy of Science in Practice: Nancy Cartwright
and the Nature of Scientific Reasoning*, ed. Hsiang-Ke Chao, Reiss Julian and
Chen Szu-Ting, New York: Springer.
- Frigg, Roman and James Nguyen. "The Turn of the Valve: Representing with
Material Models." Unpublished Manuscript.
- Frigg, Roman and Ioannis Votsis. 2011. "Everything You Always Wanted to Know
About Structural Realism but Were Afraid to Ask." *European Journal for
Philosophy of Science* 1:227–76.
- Frisch, Mathias. 2015. "Users, Structures, and Representation." *British Journal for the
Philosophy of Science* 66:285-306.
- Galton, Antony 2001. "Space, Time, and the Representation of Geographical Reality."
Topoi 20:173-87.
- Gelfert, Axel. forthcoming. "The Ontology of Models." In *Springer Handbook of
Model-Based Science*, ed. Lorenzo Magnani and Tommaso Bertolotti, Berlin
and New York: Springer.
- Giere, Ronald N. 1988. *Explaining Science: A Cognitive Approach*. Chicago:
Chicago University Press.
- Giere, Ronald N. 1994. "No Representation without Representation." *Biology and
Philosophy* 9:113 - 20.
- Giere, Ronald N. 1996. "Visual Models and Scientific Judgement " In *Picturing
Knowledge: Historical and Philosophical Problems Concerning the Use of Art
in Science*, ed. Brian S. Baigrie, 269-302. Toronto University of Toronto
Press.
- Giere, Ronald N. 2004. "How Models Are Used to Represent Reality." *Philosophy of
Science* 71:742–52.
- Giere, Ronald N. 2009. "Why Scientific Models Should Not Be Regarded as Works
of Fiction." In *Fictions in Science. Philosophical Essays on Modelling and
Idealization*, ed. Mauricio Suárez, 248-58. London: Routledge.
- Giere, Ronald N. 2010. "An Agent-Based Conception of Models and Scientific
Representation." *Synthese* 172:269 - 81.

- Glymour, Clark. 2013. "Theoretical Equivalence and the Semantic View of Theories." *Philosophy of Science* 80:286-97.
- Godfrey-Smith, Peter. 2009. "Models and Fictions in Science." *Philosophical Studies* 143:101-16.
- Godfrey-Smith, Peter. 2006. "The Strategy of Model-Based Science." *Biology and Philosophy* 21:725-40.
- Goodman, Nelson. 1972. "Seven Strictures on Similarity." In *Problems and Projects*, ed. Nelson Goodman, 437-46. Indianapolis and New York: Bobbs-Merrill.
- Goodman, Nelson. 1976. *Languages of Art*. Indianapolis and Cambridge: Hackett.
- Grüne-Yanoff, Till. 2009. "Learning from Minimal Economic Models." *Erkenntnis* 70:81 - 99.
- Grüne-Yanoff, Till and Paul Schweinzer. 2008. "The Roles of Stories in Applying Game Theory." *Journal of Economic Methodology* 15 131-46.
- Hacking, Ian. 1983. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.
- Hale, Susan. 1988. "Spacetime and the Abstract/Concrete Distinction." *Philosophical Studies* 53:85-102.
- Halvorson, Hans. 2012. "What Scientific Theories Could Not Be." *Philosophy of Science* 79:183-206.
- Halvorson, Hans. forthcoming. "Scientific Theories." In *The Oxford Handbook of Philosophy of Science*, ed. Paul Humphreys, Oxford: Oxford University Press.
- Harris, Todd. 2003. "Data Models and the Acquisition and Manipulation of Data." *Philosophy of Science* 70:1508-17.
- Hartmann, Stephan. 1995. "Models as a Tool for Theory Construction: Some Strategies of Preliminary Physics." In *Theories and Models in Scientific Processes (Poznan Studies in the Philosophy of Science and the Humanities 44)*, ed. William E. Herfel, Wladiyslaw Krajewski, Ilkka Niiniluoto and Ryszard Wojcicki, 49-67. Amsterdam and Atlanta: Rodopi.
- Hartmann, Stephan. 1999. "Models and Stories in Hadron Physics." In *Models as Mediators. Perspectives on Natural and Social Science*, ed. Mary Morgan and Margaret Morrison, 326-46. Cambridge: Cambridge University Press.
- Hellman, Geoffrey. 1989. *Mathematics without Numbers: Towards a Modal-Structural Interpretation*. Oxford: Oxford University Press.

- Hellman, Geoffrey. 1996. "Structuralism without Structures." *Philosophia Mathematica* 4:100-23.
- Hendry, Robin F. 1998. "Models and Approximations in Quantum Chemistry." In *Idealization IX: Idealization in Contemporary Physics*, ed. Niall Shanks, 123-42. Amsterdam: Rodopi.
- Hesse, Mary. 1963. *Models and Analogies in Science*. London: Sheed and Ward.
- Hodges, Wilfrid. 1997. *A Shorter Model Theory*. Cambridge: Cambridge University Press.
- Hughes, Richard I. G. 1997. "Models and Representation." *Philosophy of Science* 64:S325-S36.
- Hughes, Richard I. G. 1998. "Laws of Nature, Laws of Physics, and the Representational Account of Theories." *ProtoSociology* 12:113-43.
- Hughes, Richard I. G. . 2010. *The Theoretical Practises of Physics: Philosophical Essays*. Oxford: Oxford Univeristy Press.
- Jones, Martin. 2005. "Idealization and Abstraction: A Framework." In *Idealization XII: Correcting the Model-Idealization and Abstraction in the Sciences*, ed. Martin Jones and Nancy Cartwright, 173-218. Amsterdam: Rodopi.
- Kalderon, Mark Eli. 2005. "Fictionalism in Metaphysics " Oxford University Press
- Kennedy, Ashley Graham 2012. "A Non Representationalist View of Model Explanation." *Studies in History and Philosophy of Science* 43:326–32.
- Ketland, Jeffrey. 2004. "Empirical Adequacy and Ramsification " *The British Journal for the Philosophy of Science* 55:287-300.
- Klein, Colin. 2013. "Multiple Realizability and the Semantic View of Theories." *Philosophical Studies* 163:683-95.
- Klein, Ursula. 2001. "Tools and Modes of Representation in the Laboratory Sciences." Kluwer Academic Publishers
- Knuuttila, Tarja. 2005. *Models as Epistemic Artefacts: Toward a Non-Representationalist Account of Scientific Representation*. PhD: University of Helsinki.
- Knuuttila, Tarja. 2009. "Isolating Representations Versus Credible Constructions? Economic Modelling in Theory and Practice." *Erkenntnis* 70:59-80.
- Knuuttila, Tarja. 2011. "Modelling and Representing: An Artefactual Approach to Model-Based Representation." *Studies in History and Philosophy of Science* 42:262–71.

- Knuuttila, Tarja and Mieke Boon. 2011. "How Do Models Give Us Knowledge? The Case of Carnot's Ideal Heat Engine." *European Journal for Philosophy of Science* 1:309-34.
- Kralemann, Björn and Claas Lattmann. 2013. "Models as Icons: Modeling Models in the Semiotic Framework of Peirce's Theory of Signs." *Synthese* 190:3397-420.
- Kroes, Peter. 1989. "Structural Analogies between Physical Systems." *British Journal for the Philosophy of Science* 40:145-54.
- Kulvicki, John. 2006. "Pictorial Representation." *Philosophy Compass* 1:535-46.
- Künne, Wolfgang. 2003. *Conceptions of Truth*. Oxford: Clarendon Press.
- Landry, Elaine 2007. "Shared Structure Need Not Be Shared Set-Structure." *Synthese* 158:1-17.
- Laurence, Stephen and Eric Margolis. 1999. "Concepts and Cognitive Science." In *Concepts: Core Readings*, ed. Stephen Laurence and Eric Margolis, 3-81. Cambridge MA: MIT Press.
- Laymon, Ronald. 1990. "Computer Simulations, Idealizations and Approximations." *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 2:519-34.
- Levy, Arnon. 2012. "Models, Fictions, and Realism: Two Packages." *Philosophy of Science* 79:738-48.
- Levy, Arnon. 2015. "Modeling without Models." *Philosophical Studies* 152:781-98.
- Liu, Chuang. 1999. "Explaining the Emergence of Cooperative Phenomena." *Philosophy of Science* 66:S92-S106.
- Liu, Chuang. 2012. "A Study of Model and Representation Based on a Duhemian Thesis." In *Philosophy and Cognitive Science: Western and Eastern Studies*, ed. Lorenzo Magnani and Ping Li, 115-41. Berlin-Heidelberg: Springer-Verlag.
- Liu, Chuang. 2013. "Deflationism on Scientific Representation." In *Epsa 11 Perspectives and Foundational Problems in Philosophy of Science*, ed. Vassilios Karakostas and Dennis Dieks, 93-102. Dordrecht: Springer.
- Liu, Chuang. 2015. "Symbolic Versus Modelistic Elements in Scientific Modeling." *Theoria* 30:287-300.
- Lloyd, Elisabeth. 1984. "A Semantic Approach to the Structure of Population Genetics." *Philosophy of Science* 51:242-64.

- Lopes, Dominic. 2004. *Understanding Pictures*. Oxford: Oxford University Press.
- Lynch, Michael and Steve Woolgar. 1990. *Representation in Scientific Practice*. Cambridge MA: MIT Press.
- Machover, Moshe. 1996. *Set Theory, Logic and Their Limitations*. Cambridge: Cambridge University Press.
- MacKay, Alfred F. 1968. "Mr. Donnellan and Humpty Dumpty on Referring " *Philosophical Review* 77:197-202.
- Magnani, Lorenzo. 2012. "Scientific Models Are Not Fictions: Model-Based Science as Epistemic Warfare." In *Philosophy and Cognitive Science: Western and Eastern Studies*, ed. Lorenzo Magnani and Ping Li, 1-38. Berlin-Heidelberg: Springer-Verlag.
- Maki, Uskali. 2011. "Models and the Locus of Their Truth." *Synthese* 180:47-63.
- McAllister, James W. 1997. "Phenomena and Patterns in Data Sets." *Erkenntnis* 47:217-28.
- McCloskey, Donald, N. . 1990. "Storytelling in Economics." In *Narrative in Culture. The Uses of Storytelling in the Sciences, Philosophy, and Literature*, ed. Christopher Nash, 5-22. London: Routledge.
- Michaelson, Eliot. 2013. *This and That: A Theory of Reference for Names, Demonstratives, and Things in Between*. PhD: University of California.
- Morgan, Mary. 2001. "Models, Stories and the Economic World." *Journal of Economic Methodology* 8:361-84.
- Morgan, Mary. 2004. "Imagination and Imaging in Model Building." *Philosophy of Science* 71:753–66.
- Morgan, Mary and Tarja Knuuttila. 2012. "Models and Modelling in Economics." In *Philosophy of Economics*, ed. Uskali Mäki, 49-87. Amsterdam: Elsevier.
- Morgan, Mary and Margaret Morrison. 1999. *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge University Press
- Morreau, Michael. 2010. "It Simply Does Not Add Up: The Trouble with Overall Similarity." *Journal of Philosophy* 107:469-90.
- Morrison, Margaret. 2008. "Models as Representational Structures " In *Nancy Cartwright's Philosophy of Science*, ed. Stephan Hartmann, Carl Hoefer and Luc Bovens, 67-90. New York: Routledge

- Morrison, Margaret. 2009. "Fictions, Representations, and Reality." In *Fictions in Science: Philosophical Essays on Modeling and Idealisation*, ed. Maricio Suárez, 110-35. New York: Routledge.
- Muller, F.A. 2011. "Reflections on the Revolution at Stanford." *Synthese* 183:87-114.
- Mundy, Brent. 1986. "On the General Theory of Meaningful Representation." *Synthese* 67:391-437.
- Nguyen, James. 2016. "On the Pragmatic Equivalence between Representing Data and Phenomena " *Philosophy of Science* 83:171-91.
- Niiniluoto, Ilkka. 1988. "Analogy and Similarity in Scientific Reasoning." In *Analogical Reasoning: Perspectives of Artificial Intelligence, Cognitive Science, and Philosophy*, ed. D.H. Helman, 271-98. Dordrecht: Kluwer.
- Niven, W. D. 1965. *The Scientific Papers of James Clerk Maxwell*. Dover Publications.
- Norton, John. 2012. "Approximation and Idealization: Why the Difference Matters." *Philosophy of Science* 79:207-32.
- Odenbaugh, Jay. 2015. "Semblance or Similarity? Reflections on Simulation and Similarity." *Biology and Philosophy* 30:277-91.
- Parker, Wendy. 2015. "Getting (Even More) Serious About Similarity." *Biology and Philosophy* 30:267-76.
- Peirce, Charles. 1932. "Principles of Philosophy and Elements of Logic." In *Collected Papers of Charles Sanders Peirce, Volumes I and II: Principles of Philosophy and Elements of Logic*, ed. Charles Hartshorne and Paul Weiss, Cambridge MA: Harvard University Press.
- Perini, Laura. 2005. "The Truth in Pictures." *Philosophy of Science* 72:262-85.
- Perini, Laura. 2005. "Visual Representation and Confirmation." *Philosophy of Science* 72:913-26.
- Perini, Laura. 2010. "Scientific Representation and the Semiotics of Pictures." In *New Waves in the Philosophy of Science*, ed. P.D Magnus and J. Busch, 131-54. New York: Macmillan.
- Pero, Francesca and Mauricio Suárez. 2016. "Varieties of Misrepresentation and Homomorphism." *European Journal for Philosophy of Science* 6:71-90.
- Peschard, Isabelle. 2011. "Making Sense of Modeling: Beyond Representation " *European Journal for Philosophy of Science* 1:335-52.

- Pincock, Christopher. 2012. *Mathematics and Scientific Representation*. Oxford: Oxford University Press.
- Pissarides, Christopher A. 1992. "Loss of Skill During Unemployment and the Persistence of Unemployment Shocks." *Quarterly Journal of Economics* 107:1371-91.
- Portides, Demetris. 2005. "Scientific Models and the Semantic View of Theories." *Philosophy of Science* 72:1287-89.
- Portides, Demetris. 2010. "Why the Model-Theoretic View of Theories Does Not Adequately Depict the Methodology of Theory Application." In *Epsa Epistemology and Methodology of Science*, ed. Maricio Suárez, Mauro Dorato and Miklós Rédei, 211-20. Dordrecht: Springer.
- Portides, Demetris. forthcoming. "Models and Theories." In *Springer Handbook of Model-Based Science*, ed. Lorenzo Magnani and Tommaso Bertolotti, Berlin and New York: Springer.
- Poznic, Michael. 2015. "Representation and Similarity: Suárez on Necessary and Sufficient Conditions of Scientific Representation." *Journal for General Philosophy of Science Online* first. DOI:10.1007/s10838-015-9307-7:
- Purves, Gordon Michael. 2013. "Finding Truth in Fictions: Identifying Non-Fictions in Imaginary Cracks." *Synthese* 190:235–51.
- Putnam, Hilary. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Putnam, Hilary. 2002. *The Collapse of the Fact-Value Distinction*. Cambridge MA: Harvard University Press.
- Quine, Willard Van Orman. 1969. *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Ramsey, Jeffrey L. 2006. "Approximation." In *The Philosophy of Science: An Encyclopedia*, ed. Sahotra Sarkar and Jessica Pfeifer, 24-27. New York: Routledge.
- Redhead, Michael. 2001. "The Intelligibility of the Universe." In *Philosophy at the New Millennium*, ed. Anthony O'Hear, Cambridge: Cambridge University Press.
- Reimer, Marga and Eliot Michaelson. 2014. "Reference."
- Reiss, Julian. 2012. "The Explanation Paradox." *Journal of Economic Methodology* 19:43-62.

- Resnik, Michael D. 1997. *Mathematics as a Science of Patterns*. Oxford: Oxford University Press.
- Rickart, Charles E. 1995. *Structuralism and Structure: A Mathematical Perspective*. Singapore: World Scientific Publishing.
- Rosen, Gideon. 2014. "Abstract Objects."
- Rosenblueth, Arturo and Norbert Wiener. 1945. "The Role of Models in Science." *Philosophy of Science* 12:316-21.
- Rusanen, Anna-Mari and Otto Lappi. 2012. "An Information Semantic Account of Scientific Models." In *Epsa Philosophy of Science: Amsterdam 2009*, ed. Henk W. de Regt, Stephan Hartmann and Samir Okasha, 315-28. Dordrecht: Springer.
- Russell, Bertrand. 1919/1993. *Introduction to Mathematical Philosophy*. London and New York: Routledge.
- Salis, Fiora. 2013. "Fictional Entities."
- Schelling, Thomas. 1978. *Micromotives and Macrobehavior*. New York: Norton.
- Schier, Flint. 1986. *Deeper in Pictures: An Essay on Pictorial Representation*. Cambridge: Cambridge University Press.
- Shapiro, Stewart. 1983. "Mathematics and Reality." *Philosophy of Science* 50:523-48.
- Shapiro, Stewart. 1997. *Philosophy of Mathematics: Structure and Ontology*. Oxford: Oxford University Press.
- Shapiro, Stewart. 2000. *Thinking About Mathematics*. Oxford: Oxford University Press.
- Shech, Elay. 2014. "Scientific Misrepresentation and Guides to Ontology: The Need for Representational Code and Contents." *Synthese* 192: 3463-3485.
- Shepard, Roger N. 1980. "Multidimensional Scaling, Tree-Fitting, and Clustering " *Science* 210:390-98.
- Sismondo, Sergio and Nicholas Chrisman. 2001. "Deflationary Metaphysics and the Nature of Maps." *Philosophy of Science (Proceedings)* 68:38-49.
- Sklar, Lawrence. 2000. *Theory and Truth. Philosophical Critique within Foundational Science*. Oxford: Oxford University Press.
- Smith, Leonard A. 2007. *Chaos: A Very Short Introduction*. Oxford: Oxford University Press.
- Sterelny, Kim and Paul E. Griffiths. 1999. *Sex and Death: An Introduction to Philosophy of Biology*. Chicago and London: University of Chicago Press.

- Sterrett, Susan G. 2006. "Models of Machines and Models of Phenomena." *International Studies in the Philosophy of Science* 20:69–80.
- Stich, Stephen and Ted Warfield. 1994. "Mental Representation. A Reader." Blackwell
- Suárez, Mauricio. 2003. "Scientific Representation: Against Similarity and Isomorphism." *International Studies in the Philosophy of Science* 17:225-44.
- Suárez, Mauricio. 2004. "An Inferential Conception of Scientific Representation." *Philosophy of Science* 71:767-79.
- Suárez, Mauricio. 2009. "Scientific Fictions as Rules of Inference." In *Fictions in Science: Philosophical Essays in on Modeling and Idealization* ed. Mauricio Suárez, 158-78. New York Routledge.
- Suárez, Mauricio. 2010. "Scientific Representation." *Philosophy Compass* 5:91-101.
- Suárez, Mauricio. 2015. "Deflationary Representation, Inference, and Practice." *Studies in History and Philosophy of Science* 49:36-47.
- Suárez, Mauricio and Albert Solé. 2006. "On the Analogy between Cognitive Representation and Truth." *Theoria* 55:39-48.
- Sugden, Robert. 2000. "Credible Worlds: The Status of Theoretical Models in Economics." *Journal of Economic Methodology* 7:1–31.
- Sugden, Robert. 2009. "Credible Worlds, Capacities and Mechanisms." *Erkenntnis* 70:3-27.
- Suppes, Patrick. 1962/1969. "Models of Data." In *Studies in the Methodology and Foundations of Science: Selected Papers from 1951 to 1969*, ed. Patrick Suppes, 24-35. Dordrecht: Dordrecht.
- Suppes, Patrick. 1970. "Set-Theoretical Structures in Science." Stanford University, Stanford
- Suppes, Patrick. 2002. *Representation and Invariance of Scientific Structures*. Stanford: CSLI Publications.
- Suppes, Patrick 1960. "A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences." In *Studies in the Methodology and Foundations of Science: Selected Papers from 1951 to 1969*, ed. Patrick Suppes, 10-23. Dordrecht Reidel 1969.
- Swoyer, Chris. 1991. "Structural Representation and Surrogate Reasoning." *Synthese* 87:449-508.
- Tal, Eran. 2015. "Measurement in Science."

- Tegmark, Max. 2008. "The Mathematical Universe." *Foundations of Physics* 38: 101-50.
- Teller, Paul. 2001. "Twilight of the Perfect Model Model." *Erkenntnis* 55:393–415.
- Teller, Paul. 2001. "Whither Constructive Empiricism." *Philosophical Studies* 106:123-50.
- Thomson-Jones, Martin. 2010. "Missing Systems and Face Value Practice." *Synthese* 172:283 - 99.
- Thomson-Jones, Martin. 2011. "Structuralism About Scientific Representation." In *Scientific Structuralism*, ed. Alisa Bokulich and Peter Bokulich, 119-41. Dordrecht: Springer.
- Thomson-Jones, Martin. 2012. "Modeling without Mathematics." *Philosophy of Science* 79:761-72.
- Toon, Adam. 2010. "Models as Make-Believe." In *Beyond Mimesis and Convention: Representation in Art and Science*, ed. Roman Frigg and Matthew Hunter, 71-96. Berlin Springer.
- Toon, Adam. 2010. "The Ontology of Theoretical Modelling: Models as Make-Believe." *Synthese* 172:301–15.
- Toon, Adam. 2011. "Playing with Molecules." *Studies in History and Philosophy of Science* 42:580-89.
- Toon, Adam. 2012. *Models as Make-Believe. Imagination, Fiction and Scientific Representation*. Basingstoke: Palgrave Macmillan.
- Toon, Adam. 2012. "Similarity and Scientific Representation." *International Studies in the Philosophy of Science* 26:241-57.
- Tversky, Amos. 1977. "Features of Similarity." *Psychological Review* 84:327-52.
- Tversky, Amos and Itamar Gati. 1978. "Studies of Similarity." In *Cognition and Categorization*, ed. Eleanor Rosch and Barbara Lloyd, 79-98. Hillside New Jersey Lawrence Elbaum Associates.
- Ubbink, J.B. 1960. "Model, Description and Knowledge." *Synthese* 12:302-19.
- Vaihinger, Hans. 1911/1924. *The Philosophy of 'as If': A System of the Theoretical, Practical, and Religious Fictions of Mankind*. 1924 English Translation London: Kegan Paul.
- van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- van Fraassen, Bas C. 1981. "Theory Construction and Experiment: An Empiricist View." *Philosophy of Science (Proceedings)* Vol. 2:663-77.

- van Fraassen, Bas C. 1985. "Empiricism in the Philosophy of Science." In *Images of Science: Essays on Realism and Empiricism with a Reply from Bas C. Van Fraassen*, ed. Paul M. Churchland and Clifford A. Hooker, 245–308. Chicago and London: University of Chicago Press.
- van Fraassen, Bas C. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.
- van Fraassen, Bas C. 1991. *Quantum Mechanics: An Empiricist View*. Oxford: Oxford University Press.
- van Fraassen, Bas C. 1995. "A Philosophical Approach to Foundations of Science." *Foundations of Science* 1:5-9.
- van Fraassen, Bas C. 1997. "Structure and Perspective: Philosophical Perplexity and Paradox." In *Logic and Scientific Methods*, ed. Marisa L. Dalla Chiara, 511-30. Dordrecht: Kluwer.
- van Fraassen, Bas C. 2002. *The Empirical Stance*. New Haven and London: Yale University Press.
- van Fraassen, Bas C. 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.
- Vorms, Marion. 2011. "Representing with Imaginary Models: Formats Matter." *Studies in History and Philosophy of Science* 42:287–95.
- Vorms, Marion. 2012. "Formats of Representation in Scientific Theorising." In *Models, Simulations, and Representations*, ed. Paul Humphreys and Cyrille Imbert, New York: Routledge
- Walton, Kendal L. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Cambridge MA.: Harvard University Press.
- Warmbröd, Ken. 1992. "Primitive Representation and Misrepresentation." *Topoi* 11:89-101.
- Weisberg, Michael. 2007. "Who Is a Modeler?" *The British Journal for the Philosophy of Science* 58:207-33.
- Weisberg, Michael. 2012. "Getting Serious About Similarity." *Philosophy of Science* 79:785 - 94.
- Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.
- Weisberg, Michael. 2015. "Biology and Philosophy Symposium on Simulation and Similarity: Using Models to Understand the World: Response to Critics." *Biology and Philosophy* 30:299-310.

- Wigner, Eugene. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences." *Communications on Pure and Applied Mathematics* 13:1-14.
- Winsberg, Eric. 2009. "A Function for Fictions: Expanding the Scope of Science." In *Fictions in Science: Philosophical Essays in on Modeling and Idealization* ed. Mauricio Suárez, 179-91. New York: Routledge.
- Woods, John. 2014. "Against Fictionalism." In *Model-Based Reasoning in Science and Technology: Theoretical and Cognitive Issues*, ed. Lorenzo Magnani, 9-42. Berlin Heidelberg: Springer.
- Woodward, James. 1989. "Data and Phenomena." *Synthese* 79:393-472.
- Woody, Andrea I. 2000. "Putting Quantum Mechanics to Work in Chemistry: The Power of Diagrammatic Representation." *Philosophy of Science* 67:S612-S27.
- Woody, Andrea I. . 2004. "More Telltale Signs: What Attention to Representation Reveals About Scientific Explanation." *Philosophy of Science* 71:780 - 93.
- Yablo, Stephen. 2014. *Aboutness*. Princeton: Princeton University Press.
- Yaghmaie, Aboutorab. 2012. "Reflexive, Symmetric and Transitive Scientific Representations."