Information and Scientific Representation

Samuel D. Taylor (3887294)

S.D.Taylor@Students.uu.nl

A thesis presented for the degree of Research Master of Arts in Philosophy, August 1st, 2015.

Supervisors:

Prof. Dr. Albert Visser (UU)

&

Prof. Dr. F.A. Muller (EUR)

Department of Philosophy & Religion Studies Universiteit Utrecht, The Netherlands

"The concept of a perspicuous representation is of fundamental significance for us. It characterises the form of the account we give, the way we look at things".

Ludwig Wittgenstein, *Philosophical Investigations*: § 122

[Information is] a name for the content of what is exchanged with the outer world as we adjust to it, and make our adjustment felt upon it.

Wiener, *The Human Use of Human Beings: Cybernetics and Society*: 17

CLASSIC PEANUTS®/ by Charles Schulz

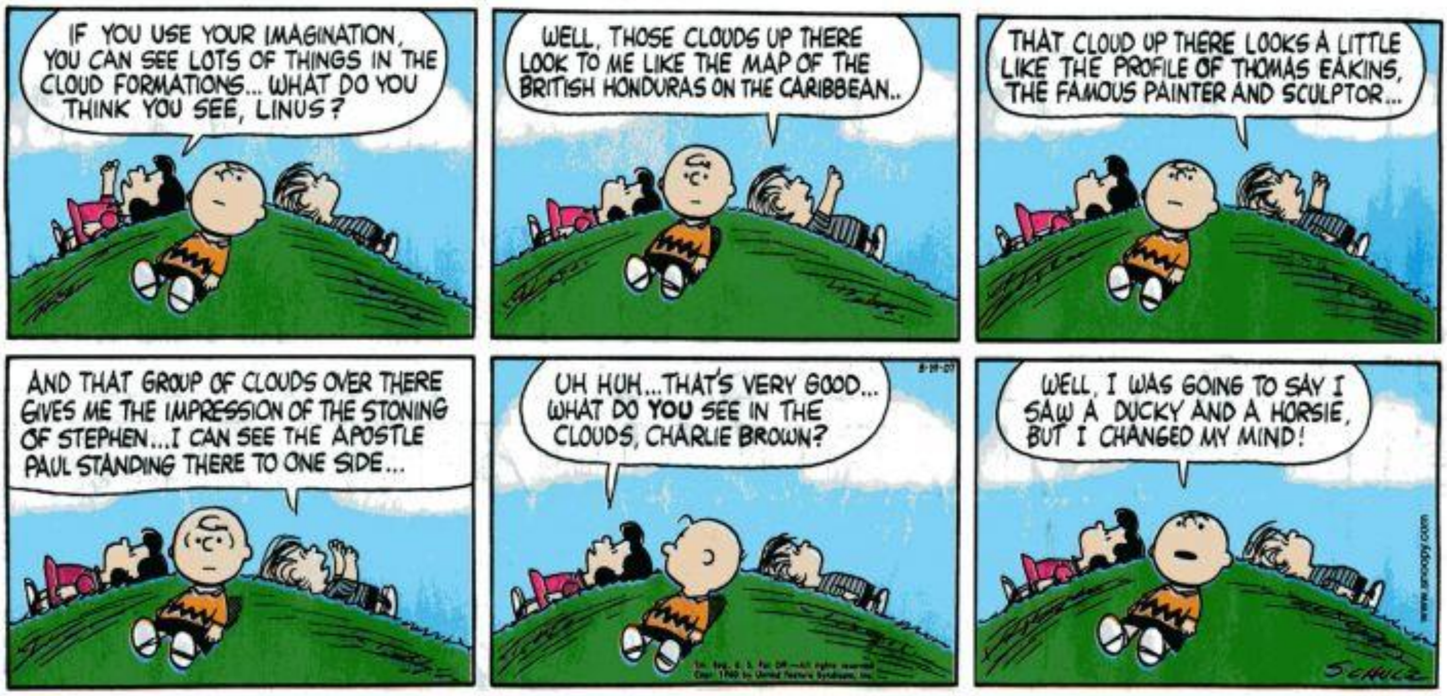

Preface

This essay discusses the way that agents come to be informed by means of scientific representations. In this respect, I will be concerned with the way that our concepts of *information* and *scientific representation* relate. This topic falls between the philosophy of science and the philosophy of information, two sub-disciplines that already exert a considerable influence on the perception, development, and practice of philosophy. In fact, one could even submit, no doubt tendentiously, that the philosophy of science and the philosophy of information are now some of the most vibrant and dynamic areas of philosophical research.

My own interest in the topics of scientific representation and information has largely been stimulated by the philosophical research community in the Netherlands. In particular, I owe a debt of gratitude to all those philosophers who believe that philosophy as a discipline should not be isolationist, but, rather, should interact with, react to, and be informed by other domains of inquiry. I have been fortunate that the prevailing sentiment in the Netherlands has been to encourage and support such interdisciplinary collaboration, and I thank all of those researchers who in their own research have made it possible to bring philosophy into contact with other areas of science.

It will be clear, then, why my attempt in this essay to think through the relationship between information and scientific representations is far from unheralded. In fact, my purpose in this essay must be situated against the backdrop of interdisciplinary currents that have been flowing now for quite some time; in particular those interdisciplinary approaches that have sought to bring philosophy into contact with science and information-theory. It is all the more surprising, therefore, that a detailed study of the relation between scientific representation and information has not already been undertaken. One reason for this neglect amongst philosophers may be that there is currently no consensus about how we should conceptualise *either* scientific representation *or* information. My ambition, therefore, is twofold. Firstly, to give a novel conceptualisation of both scientific representation and information. And, secondly, to give a novel conceptualisation of the informational nature and function of scientific representations.

It would be partly correct, therefore, to say that my interest in this essay is to survey and analyse prior philosophical engagement with scientific representation and information. But it would also be partly correct to say that my interest in this essay is to cultivate the beginnings of a novel, unified approach to philosophical engagement with scientific representations *and* information. As such, one would be right to classify this essay as both reactionary and reactive at the same time. Reactionary in the sense that I mean to prolong and fortify the philosophical clamour surrounding both scientific representation and information. But reactive in that I aim to reshape the character of this clamour by giving a novel account of the nature and function of scientific representation in terms of the nature of information. Thus, the central questions that will guide my investigation are: What conceptualisation of information and representation should we endorse? And what do these conceptualisations reveal about what scientific representations and information are and do?

The question that arises at this point is: Why? Why, that is, should one bother to cultivate an informational perspective on scientific representation in the first place? The answer I give is straightforward: because there is evidently some relation between information and scientific representation and we do not as yet have a clear conceptualisation of what that relation looks like. Our own ignorance in this regard should, I think, be justification enough for any of academic bent. But I also think that by conceptualising the relation between scientific representation and information we can attain a hitherto unavailable, informational perspective on scientific representations themselves. This essay, then, can be conceived of as the first attempt to develop a theoretical framework from which to shed light on scientific representing as an inherently informational activity. What effect this informational perspective will have on our understanding of science itself will remain open for future research. But if my discussion of information and scientific representation is even able to bring this question up for discussion, then I will be more than content.

The writing of this essay took place during the spring of 2015 as a requirement for the completion of my Research Master's degree in Philosophy at Utrecht University, The Netherlands. I am thankful to my first supervisor Prof. Dr. Albert Visser for his invaluable and profound advice, critique, and guidance throughout the writing process. I am thankful too to my second supervisor Prof. Dr. F. A. Muller for his helpful comments throughout and for encouraging me to develop and present my own ideas. I would also like to thank Dr. Jesse Mulder, Dr. Niels van Miltenburg, and PhD candidate Dawa Ometto for their willingness to engage in philosophical discussion in an amiable and good-natured manner, and for the many pieces of constructive advice they have each offered throughout my time in Utrecht. The same goes to all of my fellow cohort in the research master's program in Utrecht. More specifically, I thank Allert van Westen, Andries De Jong, Bart Kuijsten, Clint Verdonschot, Daan Dronkers, Henk van Gils, Stijn van der Leest, and most of all Ruben Noorloos, for entertaining my philosophical rants and for creating such a vibrant and inclusive research atmosphere in the philosophy department. I thank my family for their unwavering support. And, most importantly, I thank Maria Coţofan for putting up with me along the way despite my many faults – my debt to you above everyone else is beyond words.

List of Figures

Contents

Introduction

A Day at the Museum

At the NEMO science museum in Amsterdam, one finds, hidden amongst the expected array of fossils, plastinated body parts, and plasma globes, a number of exhibits designed to engender an increase in the public understanding of modern science and technology. Examples of such exhibits include: a two-meter-high plastic model of the double helix structure of DNA; a three-metre-wide foam cross section model of the structure of the human cell; an eight-metre-long model replica of the Erasmus Bridge in Rotterdam scaled down to one hundredth its actual size; and a large, multi-coloured painting of Einstein's Field Equations as taken from his General Theory of Relativity.

Each of these exhibits, like all other exhibits in the museum, is accompanied by an exhibit description – a short, simply-worded explanation of what is being exhibited and how the exhibit can gives us a greater understanding of the world. This description is intended to impart information in a way that is accessible for those who may not have previously been acquainted with the relevant theories in genetics, biology, engineering, or physics. As someone who has received some basic education in the sciences listed above, however, one finds that practically all of the information imparted by the exhibit description can be obtained by simply engaging with the exhibits themselves. That is, by studying the features of the exhibits and recalling how a particular theory in, say, biology or physics, accounts for those features. What's more, even those people who were unfamiliar with the relevant scientific theories prior to their visit to NEMO find that after reading the accompanying exhibit descriptions the exhibits themselves take on a newfound capacity to inform. One way or another, then, practically all visitors to NEMO find that upon leaving the museum both the exhibit's made out of foam, metal, and plastic, and the exhibit descriptions made out of letters, words, and sentences, impart information about aspects of reality as diverse as human cells, DNA molecules, bascule bridges, and the potential topology of space-time itself.

Now, *prima facie* we do not think of either the exhibits or the exhibit descriptions at NEMO as being identical to that which they impart information about. We would not say, for example, that the foam model of the human cell or the collection of letters, words, and sentences that constitute the exhibit description of the human cell, are actually human cells. Still, we do feel inclined to accept that upon leaving NEMO nearly all visitors would find that both the exhibit and exhibition description impart information about human cells. And, equally, that nearly all visitors would find that the other exhibits and exhibit descriptions impart information about things as diverse as DNA molecules, bascule bridge mechanisms, and the potential topology of space-time itself. But how, then, is this possible? How, in other words, do lumps of foam, plastic, metal, and glass; collections of symbols marked onto pieces of paper; or blotches of paint distributed across a canvas, come to inform us about aspects of reality that obtain at the largest and smallest scales (DNA molecules and the Universe itself), and across regions of space-time both local (spanning the Nieuw Maas in the Netherland's second largest city) and as broad as can possibly be imagined (space-time itself)?

By asking these questions we are brought to the central focus of this essay. The reason for this is because many scientists (and perhaps even ordinary layman) would be inclined to answer the questions posed above by stating that the exhibits and exhibit descriptions impart information about the various aspects of reality in virtue of being *representations* of the various aspects of reality. In fact, in virtue of being a particular sub-species of representation: the representations of science. So it follows that in the example under consideration most scientists would propose that the exhibits and exhibit descriptions at NEMO impart *information* about some aspect of reality in virtue of being *scientific representations*. But to this response the philosopher of science must query: What do we mean by *scientific representation* in this instance? And the philosopher more generally must ask: What do we mean by *information* in this instance? Without an answer to these two questions the claim that the exhibits and exhibit descriptions at NEMO impart information in virtue of being scientific representations is at best ungrounded and at worst purely speculative. And this state of affairs is even more problematic when we come to realise that the scientific sub-species of representation encompasses not only the exhibits and exhibit descriptions at NEMO, but also scientific models, scientific laws, and perhaps even scientific theories of all kinds.

Opening Remarks: Aims and Assumptions

This essay treats of scientific representation and information. It is an attempt to answer the question: How do scientific representations inform?, by attending to the questions: How should we conceptualise *scientific representation*?; and: How should we conceptualise *information*? In this respect it can be embedded within a recent trends in both the philosophy of science and the philosophy of information. My aims are as follows:

(1) To reconceptualise *representation* in such a way so as to overcome the shortcomings of competing contemporary conceptualisations;

(2) To differentiate *scientific representation* from other kinds of *representation;*

(3) To reconceptualise *information* in such a way so as to overcome the shortcomings of competing contemporary conceptualisations;

(4) To synthesise my newly reconceptualised concepts of *representation*, and *information* into one concept: *informational representation.*

(5) To differentiate *informational scientific representation* from other kinds of *informational representation;*

Before moving ahead, I think it is important to be clear on the assumptions that buttress my investigation. Firstly, I assume that:

Assumption 1: All science aims to inform us about *something*.

I take this assumption to be indisputable given our ordinary understanding of science. It simply is the case, I hold, that all science at least aims to inform us about something; whether that be the structure of financial markets, the properties of the Higgs Bosons, the interpretation of historical events, the results of chemical processes, or the taxonomical organisation of organic life.

But how then, we may ask, does science aim to inform us about such things? My answer is that science aims to inform us by means of the production of scientific representations. Thus, my second assumption is that:

Assumption 2: One major aim of science is to produce representations of the world.

My use of the term 'world' here is loose and non-committal. I certainly do mean to imply that science is in the business of producing representations of observable and actual phenomena, such as aardvark populations in Africa or star clusters in the Milky Way. But I do not necessarily restrict myself to observable or actual phenomena. It may be the case, for instance, that no one has ever observed an electron or the centre of the sun, but I do not want to suggest that science cannot, in principle, produce representations of electrons or the centre of the sun.[1]

I also hold that one major aim of science is to produce representations of the world. This may, however, be indirect in the sense that the purpose of one particular practice, say, experimentation, is only to make possible the eventual production of representations of the world. Still, I assume that the end result of science will very often be representational if that practice is to be deemed scientific at all.

---

[1] In much the same way, I do not preclude that mathematics and logic are sciences in their own right, which aim to produce representations of such things as, for instance, *abstracta* or the structure of the world/thought.

Overarching Structure

The two assumptions above form the basis of my investigation. Evidently, however, they raise a number of questions. One may ask, for example: What does it mean to say we are informed by science?; what does it mean to be informed in the first place?; what is a (scientific) representation?; how are scientific representations produced?; what is information? where does information come from?; amongst others.

These questions are difficult and can only hope to be answered given a viable conceptualisation of *(scientific) representation* and *information* as individual concepts. Giving a reconceptualisation of these concepts and considering the relation between these concepts is my primary task in this essay (aim 1, 2, 3, 4, and 5).

Firstly, I will give a reconceptualisation of *representation* (aim 1) and argue that *scientific representations* can be differentiated from other kinds of *representations* (aim 2). Then I will reconceptualise *information* (aim 3) and find that *information* and *representation* share a common conceptual form. This will set the stage for a synthesis of the concepts of *representation* and *information* in the final part of this essay. By 'synthesis' here I mean the combining of two or more things to form one unified whole. And so a conceptual synthesis is the process by which two previously distinct concepts are unified into one complex concept. In the case of *representation* and *information*, I will call the result of my conceptual synthesis: *informational representation* (aim 4). Finally, I argue that *informational scientific representation* are a special form of *informational representation*, and I give a brief account of the special nature and function of *informational scientific representations* (aim 5).

Chapter Structure

The structure of this essay is as follows. In chapter 1, I consider at length the concept of representation and discuss two competing philosophical accounts of the nature and function of representation: the *relational* view and the *functional* view.[2] The purpose of this discussion is to highlight the shortcomings of both sides of this dichotomy, and so to argue that both views fail to adequately capture what is required for a viable account of the nature and function of representation. At the close of my first chapter, I argue that any viable account of representation should accommodate both relational and functional aspects of representation. Subsequently, I argue that representation should be conceptualised as a structural relationship between three components: the representation-device (model, art-work, etc.), the thing represented, and the agent for whom the representation has a particular purpose or use. Moreover, I argue that representation is *enacted* by the relationship between these three components.

In chapter 2, I consider scientific representations in particular to illustrate how scientific representations differ from representations in other domains, such as art. I propose that the scientific production of representations is special in an important respect; namely, that it is aimed at giving the nearest approximation of reality possible. To demonstrate why scientific representations are special in this way, I consider the construction and use of the scientific representation-device: scientific

[2] In his (2009), Chakravartty divides the current views on representation between *informational* and *functional* views. As will become apparent later on, however, Chakravartty's conceptualisation of the *informational* view of representation is incommensurable with the informational account of scientific representation I develop in this essay. I will say *relational*, therefore, to avoid confusion.

models. My argument is that all models are constructed to selectively resemble their target systems, and that this is achieved by means of the constructive processes of *abstraction* and *idealisation*. Furthermore, I argue that there is a relation between the primary use of models to explain and the measurements that make model construction possible. By thinking about the construction and use of models, then, I come to the conclusion that scientific models – taken as representation-devices – *present* states, characteristics, and configurations of target systems that have been mapped or located in logical space. What's more, I argue that logical space itself is *jointly* constituted and *jointly* instantiated by agent and world. In this way, I argue that the content of scientific models is an presented relationship between agent and world. Moreover, I argue that the way that scientific models present logical space specifies how scientific representations are qualitatively different from other kinds of representations.

In chapter 3, I turn my attention to *information*. I first consider competing conceptualisations of the concept *information* as either probabilistic, algorithmic (computational), or semantic. I then argue that at its basis the difference between the various conceptualisations rests upon the pre-theoretical assumption that information must either be given a *substantive/extensionalist* conceptualisation (information 'in' world) or a *semantic/intentionalist* conceptualisation (information 'about' world; i.e. dependent upon rational agency). Then, I argue that any viable conceptualisation of information should reconcile the substantive/extensionalist and the semantic/intentionalist perspectives, and so take into account both the world and the agent, without prioritising one or the other. I propose, therefore, that we should conceptualise information as a structural relationship between three components: the information-bearer (newspaper, model, etc.), some aspect of reality, and the agent(s) for whom the information has a particular purpose or use. [3] Moreover, I argue that information is *enacted* by the relationship between these three components.

In chapter 4, I build upon chapters 1-3 to consider the relation between *information* and *(scientific) representation* as I have conceptualised them. I begin by identifying that *information* and *representation* are the same enacted structural relationship between agent, world, and artefact (i.e. model, artwork, theory etc.), and so can be synthesised into one concept: *informational representations*. I then turn to the elucidation of how informational representations are enacted. Here I consider how the artefact (i.e. the information-bearer/representation-device) makes possible the enacting of informational representations by first presenting logical space. This leads me to consider how we can know when an artefact presents logical space accurately, and hence how we can know when the artefact in question contributes to the enacting of an accurate informational representation, where the informational representation is the enacted structural relationship between agent, world, and artefact. Ultimately, I endorse a perspectival account of informational representations that suggests that informational representations capture only selected aspects of reality (i.e. the aspects mapped in logical space), and those aspects are not bits of the world in-itself, but, rather, are bits of the world as seen from the perspective of informational and

[3] The information-bearer may differ from the component aspect of reality in play (think of a newspaper reporting about the Mars Rover landing on Mars, for instance). But both are still part of the world in a broader sense.

representational human activity.[4] Finally, I return to my discussion of scientific representation to conclude that we can distinguish one kind of informational representation that is special in virtue of at least aiming to be accurate: *informational scientific representations.*

I conclude with some brief remarks about how the concept of information scientific representations could impact the philosophy of science and the philosophy of information. Most importantly, I argue that the perspectivism that is inherent to what it is to be an informational and representational being should be recognised as the central point of dispute in both the realism/anti-realism debate and the debate about the nature of information.

Key Claims

The goal of this introductory section has been twofold. Firstly, I have aimed to set out the main claims I argue for in the remainder of this essay. And, secondly, I have aimed to give a cursory outline of how my arguments will proceed. For sake of clarity, however, I think it is beneficial to provide a list of the more nuanced claims that I will defend in the course of this essay. They are as follows:

1. In order to have a clear conceptualisation of the concept *representation* we must proceed with an investigation which asks: In virtue of what do we take something to be a representation? (Chapter 1)

[4] This positions I heavily influenced by the thought of Ronald Giere (2006).

2. Both contemporary relational and contemporary functional answers to the question, in virtue of what do we take something to be a representation?, are inadequate. (Chapter 1)
3. Representation should be conceptualised as a structural relationship enacted by and between three components: the representation-device (model, art-work, etc.), the thing represented, and (the semantic activities of) an agent. (Chapter 1)
4. *Scientific representations* are special in the sense that they are taken to be aimed at giving the nearest approximation of reality possible. (Chapter 2)
5. Scientific representations make use of models that selectively resemble phenomena or data to explain. Such models are constructed by the processes of abstraction and idealisation. (Chapter 2)
6. Modelling always begin from measurement, where measurement is the location of the properties of a target system in logical space. (Chapter 2)
7. Logical space is jointly constituted and jointly instantiated by agents and world. (Chapter 2)
8. Scientific models present the content located in logical space. And therefore the content of scientific models is, in some sense, *contained* in the logical space as model construction occurs. (Chapter 2)
9. Scientific representation should be conceptualised as the structural relationship enacted by and between scientifically admissible logical spaces and a scientific model (taken as a representation-device). (Chapter 2)

10. There are many different possible conceptualisations of *information*. (Chapter 3).
11. The tension between differing conceptualisations of information is underpinned by a dichotomy between substantive/extensionalist and semantic/intentionalist conceptualisations of information. (Chapter 3)
12. To overcome the aforementioned dichotomy, information should be conceptualised as a structural relationship enacted by and between three components: the information-bearer (newspaper, model, etc.), the (semantic activities of) an agent, and the world. (Chapter 3)
13. The concepts of representation and information are, at basis, both structural relationships enacted by and between an artefact (i.e. representation-device/ information-bearer), agent, and world. And hence can be synthesised into one concept: *informational representation*. (Chapter 4)
14. The structural relationship between artefact, agent, and world is only possible when we *make sense* of how the artefact presents a logical space jointly constituted and jointly instantiated by agents and world. (Chapter 4)
15. We cannot know when our *making sense* is accurate. (Chapter 4)
16. *Informational scientific representations* are qualitatively different from other informational representations because they are at least aimed at giving the nearest approximation of reality possible – this constrains the available perspectives for making sense of informational scientific representations. (Chapter 4)

17. The world that we take ourselves to be representing or being informed about is perspectivally mediated through the relationship between agent and world (as manifest in logical space). (Conclusion)
18. The reality or truth of informational scientific representations is not to be found in a correspondence with the world or coherence of our beliefs, but rather in a shaping of the perspectival relationship between agent and world. (Conclusion)
19. The question of scientific realism versus scientific anti-realism is equivalent to the question about the nature of information. Both are questions about the nature of the perspectival relationship between agent and world. That is, about the nature of logical space (Conclusion).

Chapter 1

Representation

The purpose of this chapter is to motivate and formulate a novel conceptualisation of representation. To do this, I first consider how we should investigate the concept of representation. Then, I consider various competing conceptualisation of what representation is and is not. Before finally formulating my own conceptualisation of what representation *is*.

Representation: Methods of Investigation

What is a representation? This question requires an answer if we are to make sense of, for example, the category of 'things' to which the exhibits and exhibit descriptions at NEMO belong. Thus, if we are to make any progress on deciding what a *scientific* representation *is*, we must start by formulating a clear conceptualisation of what a representation *is*.

One option available to us at this point is to undertake a categorical investigation into what representation *is*. That is, to proceed by attempting to determine the ontological category or categories to which all representations belong. But as is clear from my examples of the exhibits and exhibit descriptions at NEMO in the introduction above, representations come in a variety of forms, including, but not limited to, abstract entities (theories, mathematical models), concrete objects (plastic models of DNA, marks of ink on an exhibit description plaque), and even processes (computer simulations, for example).[5] And so even in one domain of human activity it may be difficult to subsume all representations under one category.

Consider art, for example. What category, we may ask, will be sufficient to capture sculpture, painting, and a piece of ballet? Clearly all three can be representations in the sense of being about something, but it would be difficult to find categorical similarities between these objects that does not appeal to the trivial fact that all three are pieces of art or that all three are representations. It follows, then, that any attempt to answer the question, "what is a representation?", by means of categorisation will run into one simple and insurmountable difficulty: the only category to which all representations can be said to belong is the category of exhibiting an about-ness relation. But this categorisation is ultimately worthless, because it still leaves us without an account of what we mean by representation aside from some vague notion of about-ness. And so the ontological investigation into what a representation *is* will get us nowhere.[6]

This is important for one central reason: it illustrates that all talk of representations as being 'objects' open to categorisation is ultimately unhelpful. Why? Because a whole host of possible objects could be taken to be representations, and so we gain nothing from an object-based conceptualisation of what representation *is* without also giving an account

---

[5] One may not want to commit to the claim that representations can be processes. Regardless, this does little to undermine my point that representations are a polymorphous phenomenon.

[6] It should be noted that even if the examples of art given here are figurative, my claim is that even between so-called realist pieces of art – i.e. a portrait and a scale model of a battleship – any categorisation of different pieces of art will be ultimately based upon the about-ness relation that such objects exhibit, and so will be unhelpful.

of what makes something a representation. That is, without explaining *why* and *how* an object is a representation. And as we have seen no account of *why* and *how* an object is a representation can be given by paying attention to the categorisation of objects alone. For this reason we require another approach to the conceptualisation of what a representation *is*.

Perhaps the only other viable option – and the one expounded, to the best of my knowledge, by every philosopher currently embroiled in the debate about representation – is to approach the conceptualisation of representation by asking what we mean by the term "representation". Another way of thinking about this approach is to follow Frigg (2006) who argues that an account of what representations *are* must provide an answer to the question: 'in virtue of what [do we take] something to be a representation of something else?' (Frigg, 2006: 50). And here we seem to be on to something, because we have already noted that there appears to be a connection between about-ness and representation, but we lacked the means of unpacking this connection. So the question for us becomes: Can we attain a greater clarity about the concept of representation by pushing further with the question, "in virtue of what [do we take] something to be a representation of something else?"?

What Representation is not: Lessons from Nelson Goodman

Before moving on to consider what we gain from a Friggian-style investigation into the nature of representations (which, I think, is a lot), it will serve us well to briefly refer to the work of Nelson Goodman to specify without equivocation some of the things that representations certainly are *not*. Goodman (1976) suggests that:

> The most naïve view of representation might perhaps be put somewhat like this: "A represent B if and only if A appreciably resembles B", or "A represents B to the extent that A resembles B". Vestiges of this view, with assorted refinements, persist in most writing on representation. Yet more error could hardly be compressed into so short a formula. (Goodman, 1976: 3)

But why, then, we may ask, does Goodman think that the resemblance formula is so error-ridden? For Goodman, 'some faults are obvious' (Ibid.). For instance, the fact that resemblance, unlike representation, is reflexive, because a thing will always resemble itself to the highest degree but may not necessarily represent itself. Moreover, Goodman argues that resemblance, unlike representation, is symmetric, because for resemblance 'B is as much like A as A is like B', but the same does not necessarily hold for a representation (in Goodman's famous example, he argues that a painting of the Duke of Wellington may represent the Duke, but we would not say that the Duke himself represents the painting). Furthermore, Goodman asserts that if we conceive of representation as a form of resemblance that involves copying or imitating, then we also run into difficulties, because there can be no one correct way to represent how something *is*. As Goodman's explains with his example of the object "human being":

> [...] the object before me is a man, a swarm of atoms, a complex of cells, a fiddler, a friend, a fool, and much more. If none of these constitute the object as it is, what else might? If all are ways the object is, then none are *the* way the object is. I cannot copy all these at once (Ibid.: 7).

The resemblance account of representation, therefore, fails to specify what is being represented *without qualification* – it fails to show what the representation, taken on its own as a model exhibit, portrait etc., is supposed to resemble.[7] And this, for Goodman at least, is enough to demonstrate that representation is not simply a case of passive copying or imitating. One final Goodman quote can help to make his point here conspicuous. He says:

> The [something represented] does not sit as a docile model with its attributes neatly separated and thrust out for us to admire and portray. It is one of countless objects, and may be grouped with any of them; and for every such grouping there is an attribute of the object. (Ibid.: 32).

On this view, then, a representation can represent a man's humanity, anatomy, character, skills, or even his successes and failures. But it cannot capture what the something being represented *is* above all else. Why is this important? It is important because the investigation we will follow in the next sub-section asks, "in virtue of what [do we take] something to be a representation of something else?", and Goodman appears to have demonstrated that we should not think that something is a representation in virtue of it being a copy, imitation, or 1:1 resemblance of something else.

This negative point about what representations are not is most easily corroborated in the case of representations of human beings that, as Goodman suggests, are bodies, characters, collections of atoms etc., *all at*

[7] This qualification turns out to be central to Goodman's distinction of *representing-as*. C.f. Goodman, 1976: 27.

the same time. But in the same way we could also say, for instance, that the plastic model representation of DNA found at NEMO also fails to copy, imitate, or 1:1 resemble DNA molecules, because DNA molecules are particular sequences of nucleotides, collections of atoms, gene determinations of hair colour, and a great number of other things *all at the same time*. As both the case of a representation of a man and a representation of DNA testify, then, we should not think that something is a representation of something else in virtue of resembling, copying, or imitating that something, because a representation could never hope to resemble, copy, or imitate all of the attributes of whatever it is that is being represented.

So what is the overarching conclusion that we can draw from Goodman's critique of the common-sense conception of representation as a resembling, copying, imitating, or even duplicating of some aspect of reality? The answer, clearly, is that such a perspective on representation fails to adequately conceptualise what a representation *is*. It fails, that is, to give an adequate conceptualisation of representation. And this is exactly how we can understand Goodman's project in the opening sections of *Languages of Art*, because he endeavours throughout to determine how representation should and should not be conceptualised by evaluating differing answers to Frigg's question above. I will have need to return to Goodman briefly later on, but for now I will move on to consider in more detail the contemporary answers given to Frigg's question in the hope that a better conceptualisation of representation can be articulated.

Representation: Relational and Functional Views

Proceeding with the investigation into the concept of representation set out by Frigg above, and following on from Goodman's prescriptions as to what representations are not, we may be tempted to give a straightforward conceptualisation of a representation that goes something like this:

> A representation is something that possesses semantic properties (content, reference/denotation, truth-conditions, truth-value, etc.).

But this definition, even if correct, is not particularly helpful. Why not? Because we want to know *in virtue of what* representations possess semantic properties in the first place – put differently, we want an explanation of why representations possess semantic properties at all. It is no good, in other words, to simply say that representations have semantic properties, because we want to know *how* and *why* this is the case. To give such an explanation of the *how* and *why* it is clear that we require a deeper story about the nature and function of representations. We require a deeper story, then, about what it is that representations *are*.

To begin to give this deeper story it is first important to fix our terminology to make clear what occurs when we say that we have a case of representation. Taking our cue from Mauricio Suarez (as is now standard in philosophical discussion of representation), we can proceed as follows:

> Let us refer to the vehicle of the representation as the 'source'; and the object as its 'target'. (Thus in a portrait, the canvas is the source and the person portrayed is the target). Anything can in principle play the role of sources or targets, so these terms are mere place-

> holders. I shall assume that *X* is the source and *Y* is the target if and only if '*X* represents *Y*' is true. (Suarez, 2010: 91).[8]

If we accept that anything could in principle play the role of either a source or target, then it is evident that representations can be found in many (if not all) areas of human society. In art, for instance, we find representations of people, places, ideas, emotions etc.; just as in science we find representations of a particular features of the world (DNA molecules, human cells, bascule bridges, and the potential topology of space-time). By thinking about representation in terms of sources and targets, therefore, we have a straightforward way to consider representation more abstractly; i.e. without recourse to any particular representation in any particular context.

So where does this abstract consideration of representations lead? At the very least, Suarez's analysis of representations in terms of sources and targets illustrates that for every example of representation one can think of, it is always the case that there are at least two 'place-holders' involved, because it is always the case that at least one thing (the source) represents another (the target). It quite simply has to be this way, for it is impossible for something to be both source and target without representing itself by being itself, and then we do not appear to have an interesting case of representation at all.[9] In asking the question, what *are*

---

[8] It could be argued that the fact that representation involves a representans and a representandum is obvious. But I will make reference to Suarez to embed myself within the contemporary philosophical literature.

[9] I say that when something represents itself it is uninteresting for the reasons that a representation of this form appears to tell us nothing more than what we can already know from interacting with the 'something' in question. One argument against this view is to say that I can at time t_1 represent myself as being scared at time t, but this view trades on an endurantist view of persistence and identity that I do not feel obliged

representations?, then, it is easy to suppose we are asking for an account of the way in which a source – for example, the canvass of a portrait or the symbols constituting a mathematical equation – *relates* to a target – for example, a person or the speed of light in a vacuum. That is, that we are asking for an account of the nature and function of a representation in terms of the relation between source and target.

According to Suarez (2010), in fact, as a rule analytic inquiry into representation 'tends to presuppose that representation is a relation'. In fact, he proclaims that:

> The analytical inquiry […] in its most basic form, takes it that representation is a relation *R* such that the assertion that '*X* represents *Y*' is equivalent to the assertion that '*R* holds between *X* and *Y*' (Suarez, 2010: 91).

On this view, *X* would be a source and *Y* would be a target, with representation being the relation *R* between the two. I will say that those philosophers who offer conceptualisations of representation in terms of some relation *R* adopt a *relational view* of representation.

We will return to the various attempts to spell out the *relational view* of representation shortly in the next sub-section below. For now, however, we must briefly explain the most fundamental reason that *relational view* of representation has not been universally accepted by philosophers. The reason, to get straight to the point, is that the *relational views* of representation say nothing about that which causes a

---

to entertain here. If, therefore, we say that all regions of space-time are different, then it becomes clear why to say that a something that represents itself in space *s* at time *t* is to say nothing more than there is something in space *s* at time *t*.

representational relation to obtain. Another way of putting this is to say that the *relational views* of representation says nothing about the purpose of representations for human agents – they offer no account of *why*, for us, a representational relation would occur in the first place. As was made apparent in my earlier examples about exhibits and exhibit descriptions at NEMO, it is clear that representations are closely interconnected with the activity of human agents. And according to those philosophers who are critical of the *relational view* of representation, this can only be accounted for by prioritising the *use*, or, again, the *function*, of representations in the context of human agency and cognition. So those opposed to a relational conceptualisation of representations will, as Chakravartty (2009) put it, prioritise the 'uses [of representations] in cognitive activities performed by human agents in connection with their targets' (Chakravartty, 2009: 3). They will, in other words, attempt to formulate what I will call a non-relation, *functional view* of representation.[10]

The Relational View of Representation

I will now move ahead with Friggian-style investigation into representation currently underway. Accordingly, I will now consider the most prominent *relational views* of representation, before moving on to consider the shortcomings of such *relational views* of representation in the next sub-section.

[10] In mathematics functions are also considered to be relations of a sort, but my use of the term functions here is in the ordinary language sense of being useful, practical, handy, or purposeful. In this way, I follow Chakravartty (2009).

The *relational view* of representation 'admits of several species […] including relations of isomorphism, partial isomorphism, and homomorphism' (Chakravartty, 2009: 2). In his 2009, Chakravartty argues that all of these relational conceptualisations of representation are forms of one broader approach to the conceptualisation of representation: an approach that conceptualises representation in terms of a 'general relation of similarity' (Ibid.). For my part, I do not think that it is necessary to classify all *relational views* of representation in terms of similarity, but I do hold that all *relational views* of representation conceptualise representations as a relation of one kind or another.[11] To see why, I will now consider different relational views of representation in turn.

I will start with the conceptualisation of representation as an isomorphic relation given first by Van Fraassen in his (1980) and (1989), and as partial isomorphic relation given by French (2003) and da Costa and French (2003).[12] In a now (in)famous passage, Van Fraassen (1980) declared that:

> To present a theory is to specify a family of structures, its models; and secondly, to specify certain parts of those models (the empirical substructures) as candidates for the direct representation of observable phenomena. The structures which can be described in

[11] C.f. Bartels, 2006: 14, and Suarez and Sole, 2006: 44, for criticism of Chakravartty's claim that similarity is the same kind of relation as homomorphism and isomorphism respectively.

[12] C.f. Van Fraassen, 1980: chapter 3, 1989: chapter 9 for a more detailed exposition than I give here. It should also be noted that Van Fraassen's view on scientific representation has turned out to be more nuanced than one may originally have thought. In fact, his current account would not fit with the *relational view* at all, and is much more suited to the *functional view* I will discuss below (c.f. Van Fraassen, 2008). I will engage directly with Van Fraassen's (2008) account in the next chapter.

> experimental and measurement reports we can call *appearances*: the theory is empirically adequate if it has some model such that all appearances are isomorphic to empirical substructures of that model. (Van Fraassen, 1980: 64)

Thus a model (the substructures of which we can read in this instance as Suarez's 'source', or 'vehicle', of a representation) will be empirically adequate if 'all the actual phenomena fit inside' so that 'what it says about the observable things and events in this world is true—exactly if it 'saves the phenomena" (Ibid.: 12). And the way that phenomena 'fit inside' the source of a representation (a model) is by means of an isomorphic relation, whereby we have:

> […] a total identity of structure and [a] limiting case of embeddability: if two structures are isomorphic then each can be embedded in the other. (Ibid.: 43)

On this view, then, representation should be conceptualised as the isomorphic relation between two structures: a target (for van Fraassen, 'observable things and events in this world') and a source (for van Fraassen, the empirical substructures of a model). Thus, anything that does not satisfy this isomorphic relation between observable structures and the substructures of a model does not count as a representation. As a consequence of this stricture, in fact, Van Fraassen argues that certain models and theories that are not isomorphic to the structure of observable phenomena cannot be considered representations at all (Ibid.: Ch. 6). Van Fraassen's relational conceptualisation of representation, therefore, turns out to be a form of radical empiricist isomorphism.

A more recent partial isomorphic *relational view* of representation has been defended by French (2003). He argues, in accordance with Budd (1993), that 'a [source] represents by sharing properties with its [target]' (French, 2003: 1475), and that:

> This "sharing" of properties can then be captured via the notion of isomorphism, where this is taken to hold between the [source…] and the […target]. (Ibid.)

French advocates a *partial* isomorphic conceptualisation of representation because representations are 'regarded not to be true in the correspondence sense, but as partially true or approximately true, or as containing an element of truth' (da Costa and French, 2003: 4). Moreover, French argues that this offers a 'new perspective on the nature of representations themselves' (Ibid.).

So what is the 'new perspective' on representations that comes with the *partial* isomorphic conceptualisation? The central idea is that 'representations of the world that are not perfect copies but are, in certain respects, incomplete and partial' (Ibid.). The structure of representations, da Costa and French propose, is open, because it is a 'partial structure'. Now, at this point it will be expedient to quote French at length to see why partial structures still support an isomorphic – and hence *relational* – conceptualisation of representation. French says:

> […] a partial structure is a set-theoretic construct $a = \langle D, R_i \rangle_{i' \in \mathrm{I}}$, where D is a nonempty set and each R_i is a partial relation. A partial relation $R_i, i' \in \mathrm{I}$ over D is a relation which is not necessarily defined for all n-tuples of elements of D (see da Costa and French 1990,

> 255). Each partial relation R can be viewed as an ordered triple $\langle R_1, R_2, R_3 \rangle$ where R_1, R_2 and R_3 are mutually disjoint sets, with $R_1 \cup R_2 \cup R_3 = A^n$, and such that: R_1 is the set of n-tuples that belong to R, R_2 is the set of n-tuples that do not belong to R, and R_3 is the set of n-tuples for which it is not defined whether they belong or not to R. If we have two partial structures
>
> $$a = \langle D, R_k \rangle_{k \in \mathrm{K}}$$
>
> and
>
> $$a' = \langle D', R'_k \rangle_{k \in \mathrm{K}}$$
>
> (where R_k and R'_k are partial relations as above, so that $R_k = \langle R_{k1}, R_{k2}, R_{k3} \rangle$ and $R'_k = \langle R'_{k1}, R'_{k2}, R'_{k3} \rangle$ then a function f from D to D is a partial isomorphism between a and a'if 1) f is a bijective and 2) for all x and y in D, R_{k1xy} iff $R'_{k1}f_{(x)}f_{(y)}$ and R_{k2xy} iff $R'_{k2}f_{(x)}f_{(y)}$ (French and Ladyman, 1999; Bueno, 1997). Of course, if $R_{k3} = R'_{k3} = \emptyset$, so that we no longer have partial structures but "total" ones, then we recover the standard notion of isomorphism (see Bueno, 1997).[13] (French, 2003: 1480).

All this is to say that for French representation should be conceptualised as a partial or 'total' isomorphic relation between at least two partial structures. It makes little difference to my investigation what exactly the partial structures *are* in themselves (although French proposes that partial structures can be models and perhaps even aspects of reality itself (da Costa and French, 2003: Ch. 3)). Rather, what matters for my investigation

[13] As is evident from the square root symbols, with Da Costa and French we leave classical logic behind.

is that for French representation *is* the partial isomorphic relation between at least two partial structures; nothing more and nothing less. Thus, French, like Van Fraassen before him, conceptualises representation as an (partial) isomorphic relation.

We have now seen two species of the *relational view* of representations: isomorphism and partial isomorphism. The final *relational view* I will present is the homomorphic conceptualisation of representation. The homomorphic conceptualisation of representation holds that 'Representations should be modelled by means of faithful mappings', whereby, when we have '*A* (the domain to be represented) and *B* (the domain representing *A*)':

> a mapping $f: A \rightarrow B$ can be defined which maps A *onto* B. The mapping f is not necessarily one-to-one and satisfies two following conditions:
>
> (i) For all j and all elements a_i of A: if $R_j^B\ (f(a_1), \dots, f(a_n))$, then $R_j^A\ (a_1, \dots, a_n)$
>
> […and]
>
> (ii) For all j and all elements a_i of A: if $R_j^A\ (a_1, \dots, a_n)$, then $R_j^B\ (f(a_1), \dots, f(a_n))$,
>
> If (i) and (ii) are fulfilled, f is a homomorphism from A onto B, and B, by virtue of the existence of f, can be said to be a *homomorphic image* of A (Dunn and Hardegree 2001, 15). The structural concept of representation claims that B represents A only if B is a homomorphic image of A. (Bartels, 2006: 8).

Bartels 'structural concept of representation' should be read as a homomorphic conceptualisation of representation, because 'B represents A' is to be explained by the relation of A being homomorphic to B' (Ibid.: 9).[14]

So why is Bartel's homomorphic conceptualisation of representation a *relational view* of representation? We find the answer explicitly in his writings where he states:

> a homomorphism between *relational* structures A and B either exists or does not exist; in the first case, B represents A, whereas in the second case B does *not* represent A. (Bartel, 2006: 14, my italics)

We have now seen three variants of the *relational view* of representation: isomorphism, partial isomorphism, and homomorphism.[15] It is now time, therefore, to make explicit the features that all relation views share; namely, the commitment to a conceptualisation of representation in terms of some relation R that holds between a source and a target (see fig. 1). All *relational views* of representation, then, hold that 'the assertion that 'X represents Y' is equivalent to the assertion that 'R holds between X and Y' (Suarez, 2010: 91), where X and Y are sources

---

[14] It is important to note that Bartels does make a distinction between what he calls the 'content' of the representation and the 'reference' of the representation, where the target, or reference, is something 'intentionally' 'selected' to be represented. He argues that 'The reference of B is fixed by a representational mechanism, i.e. either by an intentional or a causal process. In contrast, the representational *content* of B is determined by B's structural properties, i.e. by the relational structure B is endowed with. But still, for Bartels at least, in order to count as a representation B must have a structural content 'generated by means of homomorphisms'. (C.f. Bartels, 2006: section 3).

[15] For other conceptualisation of representation that fit the *relational view* as I have defined it, see Pincock's (2011) account of structural isomorphism and Weisberg's (2012) account of similarity.

and targets parsed in terms of, for instance, the structure of (observable) things and events and (partial) (sub)structures of models. The basic idea, therefore, is that with the *relational view* any morphism will do to conceptualise what representation *is*. In this way, *relational views* of representation are *reductive* in that they take representation to be conceptualisable solely in the relational terms given above.

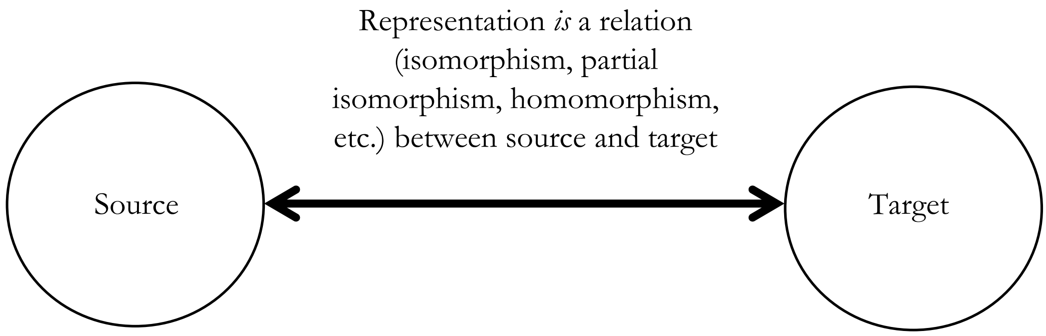

Figure 1.

Shortcomings of the Relational View of Representation

The first potential problem with the *relational view* of representation is that the obtaining of a relation between a source and a target may not even be necessary for there to be a representation. This non-necessity argument against the *relational view* of representations holds that in some cases where we seem to have a representation we do not seem to have a relation of any kind (let alone of isomorphism, partial isomorphism, or homomorphism!). To appreciate the objection on the table here, Chakravartty (2009) asks us to consider the 'the one case that does seem to present a *prima facie* difficulty for an analysis of scientific representation in terms of [a relation]: linguistic representation' (Chakravartty, 2009: 5). The idea here being that one cannot find a relation between, for example, the word "human cell" and the microscopic composite parts of our body. And thus that linguistic representation is a form of representation that is not relational at all.

I agree with Chakravartty, however, that this objection is not particular strong. The reason being that most philosophers who defend a *relational view* of representation will simply claim that we are confusing the linguistic device with the semantic content of a linguistic entity, say, the words "human cell". A linguistic device may be blotches of paint on a canvas, marks of chalk on a blackboard, or trails of ink on paper intricately formed to display the ordered symbols "human cell". But the defender of the *relational view* of representation will argue that these linguistic devices cannot be equated with their semantic content of the linguistic entity "human cell", which is instead a specific family of models that captures what we mean by the microscopic composite parts of our body.[16] Chakravartty borrows Suarez's (2003) example to emphasise the same point:

> the quantum state diffusion equation for a particle subject to a localization measurement, as written in a textbook or on a blackboard, does not appear similar to the properties of any particle, to be sure. But this, surely, is to see the equation in a

[16] The informed and attentive reader will have noted that the point here directly corresponds to the debate between advocates of the semantic view of scientific theories and advocates of the syntactic view of scientific theories. I will not discuss this debate in detail, but for an account of the syntactic view see Braithwaite, 1953; Nagel, 1961; and Spector 1965; and for an account of the semantic view see Van Fraassen, 1980; Giere, 1991; and Suppes, 1988 and 2002.

> superficial way, as merely blotches of ink on a page, or intricately-shaped trails of chalk dust. Anyone who sees only this when viewing such a thing would presumably have no grounds for thinking that it was a representation at all. On the other hand, having learned the languages of mathematics and physics, one may view the content of these blotches and trails as trajectories in phase space, and these most certainly have features in common with the states of particles subject to localization measurements. (Ibid.: 6)

The argument, then, is simple: the *content* of a linguistic representation can be given a relational conceptualisation insofar as one has grasped the semantics of the language employed. Since the same argument can be employed for all other cases in which we *prima facie* seem to have a representation without a relation, the non-necessity argument against the *relational view* of representation fails.[17]

A more troubling objection for the *relational view* of representation, however, is the non-sufficiency argument. The non-sufficiency argument holds that relations are not sufficient to conceptualise representations. Chakravartty summarises the argument as follows:

> representation is achieved only in circumstances in which agents know or have otherwise mastered the system of representation being used […] That is why relations such as similarity cannot do the job on their own; these relations only serve the goal of

[17] Chakravartty offers another non-necessity argument against the *similarity view* of representation he endorses, but since this arguments turns on the need for different kinds of *relations* for different kinds of representation, it will not undermine a view of how to conceptualise representations that does not confine itself to similarity, but focuses on relations of any kind; i.e. the *relational view* I am considering here.

> representation subject to the internalization of the semantics of their forms of expression and relevant representational conventions by their users, either by means of hard-wired cognitive responses, scientific or other training, or both. (Ibid.: 9).

The point here, then, is that a relational conceptualisation of representation cannot 'do the job on its own', because even if representation does involve a relation of one kind or another, a conceptualisation of representation must also pay attention to the *use* or *function* of a representation by an intentional agent. So a relational conceptualisation of representation will not be sufficient to capture all of what is required from a conceptualisation of representation.

In a now widely cited example, French (2003) imagines a scenario in which the wind and sea carve the Lorentz transformation into the sand on a beach. French implies that our intuition is still to think of the sand carvings as representations of 'relativistic phenomena of some sort' even without the relevant agential intention associated with the *use* or *function* of the representation (French, 2003: 1473). Thus, French wants to make space to defend that claim that a *relational view* of representation really could be sufficient to conceptualise representation. I do not share French's intuition that the sand-carving could still be a representation without the relevant agential intention associated with the *use* or *function* of the representation, and I am not alone (c.f. Suarez, 2003; Elgin, 2009; Chakravartty, 2009). The reason for the rejection of French's intuition is because it seems to me (and others) that even if the representational status of the sand-carving is independent of how the representation is constructed, this does not imply that the sand-carving can be a representation without having some *function* or *use*. Put differently, the intuition is that representations are representations for some purpose, and this purpose can only be factored in by accommodating intentionality, or agency, in one's conceptualisation of representation.[18]

One open possibility for advocates of the *relational view* of representation would be to defend a naturalistic account of intentionality (French, 2003; Bartels, 2006). But it remains to be seen if such naturalistic intentionality could avoid the charge of being parasitic on the intentions of human agents, especially in representational domains such as art or science.[19] So, without a viable account of the naturalistic source of the intentional *use* or *function* of representation, those who endorse the *relational view* face two options: (1) to conceptualise representation solely in terms of relations and ignore the troubles the *relational view* has in accommodating the *use* or *function* of a representation; or (2) to accept that the *relational view* of representations does not give sufficient conditions for representation.

The Functional View of Representation

The very possibility of (2) above has led some philosophers to reject the project of conceptualising representations as relational. These philosophers have developed their own distinct Friggian-style investigation into how best to conceptualise representation. Earlier, I called these relational heretics advocates of a *functional view* of

---

[18] Chakravartty argues in a similar vein that 'the *least* controversial feature of scientific representation is the idea of intentionality' (Chakravartty, 2009: 10).

[19] This seems like a very long shot to me, but *relational view* philosophers are welcome to try to give such an account of agent-independent intention.

representations. I will now consider the most prominent *functional views* of representation, before moving on to consider the shortcomings of such *functional views* of representation in the next sub-section.

Much like the *relational view* of representation, the *functional view* of representation admits of several species including demonstration and interpretation, exemplification, or inferential. What unites all *functional views* of representation, however, is a focus on the way that representations are used – their purpose – in regards to the tasks, both physical and cognitive, of human agents. All *functional views* of representation, therefore, emphasise the mind-dependence of representations and conceptualise representations accordingly. This is in stark contrast to most – if not all – *relational views* of representation, which, as we have seen, adopt a mind-independent, reductionist approach to the conceptualisation of representations. *Functional views* of representation, then:

> focus primarily on the question 'what is [a] representation?', where 'representation' is conceived as a set of knowledge-exercising practices, constituted by whatever it is that [we] do when engaged in the process of representing things. (Chakravartty, 2009: 13).

The first *functional view* of representation I will consider is the conceptualisation of representation as Denotation Demonstration Interpretation (DDI) given by R.I.G Hughes. Hughes, as the reader can probably guess, can easily be embedded within the tradition inaugurated by Nelson Goodman (1976).[20] Suarez (2015) describes the DDI conceptualisation of representation as follows:

> According to [Hughes'] account a source A represents a target B when the following three conditions are met: i) The source stands for the target in the sense that it denotes it; ii) Some demonstration is carried out by an agent on the model; and iii) The results of this demonstration are then interpreted, so as to apply them to the target. (Suarez, 2015: 13)

Although important for situating DDI within the larger debate about representation, we need not trouble ourselves with questions about what is meant exactly by denotation, demonstration, or interpretation here (Suarez has a nice example of this concerning the representations produced by Galileo in his (2015: 14)). Instead, it is enough for us to note that Hughes conceptualises representation as 'a set of speech acts', rather than as a straightforward relation between sources and targets (Hughes, 1996: 329).

That the concepts of denotation and interpretation employed by Hughes seem to suggest a relation of some kind, then, is not important.[21] What is important, however, is Hughes attempts to give a conceptualisation that prioritises the way that representations are used.

---

[20] An exemplary discussion (and critique) of the development of philosophical accounts of representation can be found in Muller (2010: especially 410-413).

[21] Suarez (2015) has convincingly argued that even if Hughes account of DDI is 'a hybrid of a relation (denotation, mapping), and a number of activities (demonstrating, ascribing, partitioning)', it should still be taken as at least an attempt to switch the attention away from a conceptualisation of representation in terms of relations, towards a conceptualisation of representation in terms of our practices (c.f. Suarez, 2015: 18-25).

That is, the *function* that representations have for intentional agents. Hughes' project, therefore, can be viewed as an attempt to conceptualise representation not as a relation between (sub)structures of models, on the one hand, and structures of facts, states, effects, phenomena, etc., on the other. But, instead, as something that is essentially *functional*; that is, as something that we must conceptualise first and foremost as being of use.

In much the same way as Hughes, Elgin's (2004) account of representation as exemplification must also be said to have been inspired by Nelson Goodman (1976). For Elgin, exemplification is:

> the device by which samples and examples [i.e. representations] highlight, exhibit, display, or otherwise make manifest some of their features (Elgin, 2004: 124).

Thus, exemplification depends upon the cognitive activities of agents to which representations 'make manifest' their features. In this way, 'representation affords epistemic access to features of the object that are otherwise difficult or impossible to discern' (Elgin, 2009: 80). And this, for Elgin, is the core of representation.

To see why this is the case, consider this passage from Elgin (2009: 81):

> An exemplar can exemplify only some of its properties. It brings those properties to the fore by marginalizing, downplaying, or ignoring other properties it instantiates. It may exemplify a cluster of properties, as for example a fabric swatch exemplifies its colors, texture, pattern and weave. But it cannot exemplify all of its properties. Moreover, an exemplar is selective in how precisely it

> exemplifies. A single splotch of color that instantiates dusky rose, rose, and pink may exemplify any of these properties without exemplifying the others. Although the color properties it instantiates are nested, it does not exemplify every property in the nest. Exemplars are symbols that require interpretation.

The point this passage draws out is as follows. For Elgin, we should conceptualise representations as those thing that make possible the interpretation of a certain class of exemplars; namely, the sources or vehicles of representation (portraits, models, etc.). Representation, then, on Elgin's view, is not to be conceptualised above all else as a relation between a source and a target, but as something that exemplifies certain features which facilitates certain cognitive activities of human agents. The cognitive activities that Elgin mentions include interpretation and belief-formation, but there is no doubt that other cognitive activities could also be relevant here. In short, Elgin suggests that representation 'affords resources for thinking', in one way or another (Elgin, 2009: 85). And this conceptualisation obviously promotes a *functional*, as opposed to *relational*, view of representation.

We have now considered two species of the *functional view* of representation: demonstrating and interpreting and exemplification. To round-off this sub-section, I will consider one final *functional view* of representation: the inferential conception of representation offered by Suarez (2004, 2009, 2015).[22]

[22] It should be noted that the conception of laws as inference tickets goes back to Prior and Geach in the 1960's. My focus on Suarez here, however, is directly in the context of the contemporary discussion of scientific representation. But I do want to stress that I do think that the two philosophical discussions are intertwined (c.f. Prior, 1971).

Suarez's inferential conception is concerned with the way that representations are used. This is evident from his account of 'representational force', which he argues is 'the capacity a source has to lead a competent and informed user to a consideration of a target' (Suarez, 2004: 768). In Suarez's example, a spiral staircase may be taken to 'represent the mechanics of a spring, or the structure of DNA', just in the same way that the exhibits and exhibit descriptions at NEMO in the example we started with may also be taken to represent human cells, the structure of DNA, Erasmus bridge, and the topology of space-time. Importantly, however, Suarez argues that 'the source's force varies with intended use', so that 'if an agent is competent and informed, or ambivalent, about the use of both representations, the force of the source will be double or ambiguous respectively' (Ibid.).

Now, Suarez choses to introduce and elucidate the notion of representational force because it serves to buttress his inferential conception of representation. This is the case, because Suarez is able to defend the claim that the best we can hope to do in regards to a conceptualisation of representation is 'to describe its most general features'; with representational force being perhaps *the* key features in need of description (Ibid.: 771). Suarez, therefore, makes room for 'a deflationary or minimalist attitude and strategy towards the concept of […] representation' (Ibid.).

Suarez formulates the inferential conception of representation as follows:

> *A* represents *B* only if (i) the representational force of *A* points towards *B*, and (ii) *A* allows competent and informed agents to draw specific inferences regarding *B*. (Ibid.: 773).

And in keeping with his minimalist approach to the concept of representation, he explains that:

> Part (i) leaves open the issue of how many agents are required in a […] community to fix the representational force of a source, and what the structure of the community and its practices ought to be in order to determine this force. Part (ii) leaves open the issue of what *A*'s internal structure ought to be like in order to yield correct inferences about *B*. In particular it does not require that *A* allow deductive reasoning and inference; any type of reasoning—inductive, analogical, abductive—is in principle allowed, and *A* may be anything as long as it is the vehicle of the reasoning that leads an agent to draw inferences regarding *B*.

Thus, Suarez's inferentialist conception of representation is weak enough to accommodate stronger conditions on representation as a concept, even to the point of relations of one kind or another. But, crucially, the inferentialist conceptualisation provides a way to think about representation without recourse to any stronger conditions, and so no devolution into the *relational view* of representation is permitted. All that is required on the inferentialist conception of representation, then, is that representations are conceived of as being those things that enable a representational source to 'license inferences regarding its target' (Ibid.: 776). And as Suarez explains,

> The representation is true if it licenses no inferences to false conclusions about the target; it is complete if it is true and fully informative, licensing inferences to all truths about the target; and it is empirically adequate if it is complete with respect to all the observable or measurable aspects of the target, licensing inferences to all the truths about those aspects. (Ibid.)

It should be clear, therefore, why the inferentialist conceptualisation of representation fits squarely into the *functional view* of representation. But in case it is not clear, I need only say that the inferential conception of representation aims chiefly to conceptualise representations in terms of our use of representations to make inferences. As Suarez himself says:

> the inferential conception [of] representation […] requires the correct application of *functional* cognitive powers (valid reasoning) by means that are objectively appropriate for the tasks at hand (i.e., by models that are inferentially suited to their targets). (Ibid.: 778, my italics)

We have now seen three species of the *functional view* of representation: demonstration and interpretation, exemplification, and inferential. It is now time, therefore, to make explicit – in case it is not already clear – the features that all functional views share; namely, the commitment to a conceptualisation of representation in terms of the way that agents use the representations in connection with their targets. All *functional views* of representation, then, hold that in order to give a viable conceptualisation of representation we must be clear on what function the representation has for us. We must, in other words, have an account of representations parsed in terms of what we do with representations (see fig. 2). Given these commitments, it is unsurprising that *functional views* of representation typically oppose the reduction of representation to any mind-independent set of conditions, such as a relation of one kind or another.[23]

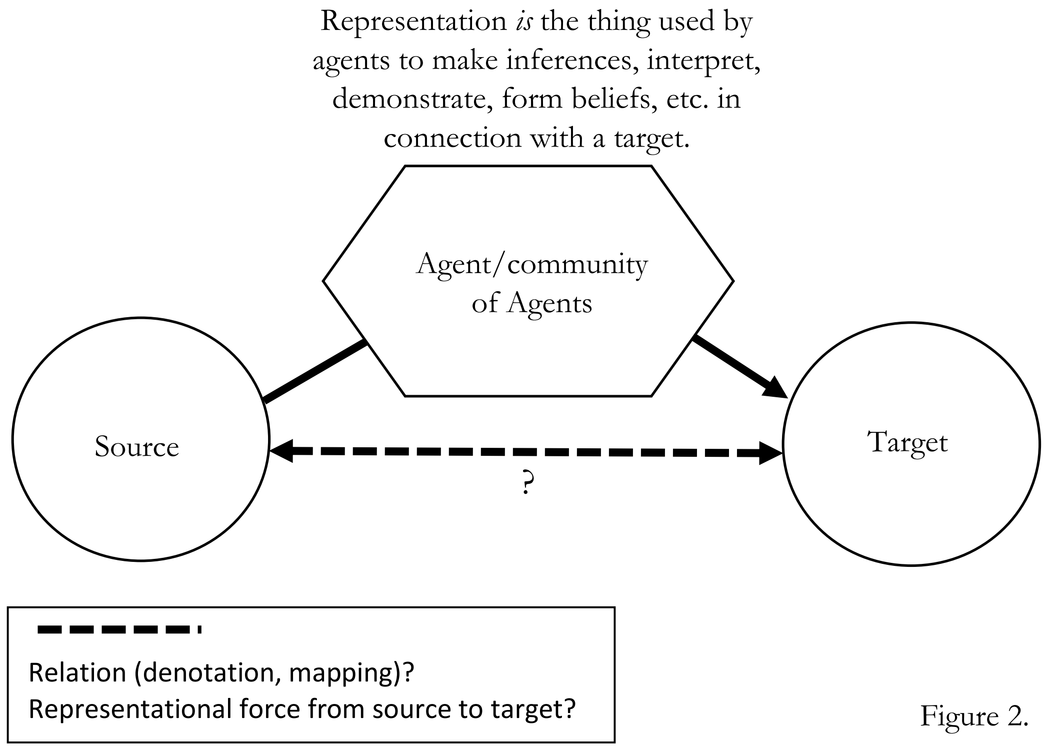

Figure 2.

Shortcomings of the Functional View of Representation

Advocates of a *functional view* of representation will always hold that a viable conceptualisation of representation must include an account of the use or

[23] For other conceptualisation of representation that fit the *functional view* as I have defined it, see Contessa's (2007x2) account regarding inference and van Frassen's (2008) account regarding use. In his (2008: 23), Van Fraassen presents the following '*haupstaz*':

> There is no representation except in the sense that some things are used, made, or taken, to represent some other things as thus or so.

function of a representation as a necessary condition. They will defend this positions by reference to the non-sufficiency argument against the *relational view* of representation discussed above. Thus, advocates of the *functional view* of representations will argue that any conceptualisation of representation that does not include functions will be incomplete, *ergo* one's conceptualisation must accommodate the uses of representations in the cognitive tasks of agents.

The central problem with the *functional view* of representation, however, is that it all but eradicates the possibility of specifying in detail the sufficient conditions of representation. This is the case, because even if we can formulate some necessary condition based upon the use or function of representations for agents, we have no way to decide which of the various different kinds of uses agents may have for representations constitute the sufficient conditions (demonstration, interpretation, inferential). Moreover, the *functional view* may not be able to specify any sufficient conditions that must be met in order for one to be sure that there is a valid case of representation in the first place. The reason is simple: the *functional view* of representation appeals only to the uses of representation and this will not be sufficient to mark the difference between accurate and false representation.

Consider, for example, the phlogiston model of heat given by physicists in the nineteenth century. In this instance, there were a great many 'representations' of phlogiston produced. These so-called representations purported to represent the properties of phlogiston, such as phlogiston's facilitation of the combustion of combustible bodies. But the phlogiston model of heat was ultimately determined to be faulty and so became obsolete: it was decided that phlogiston does not exist. The question, then, is what are we to say about the 'representations' of phlogiston produced by nineteenth century physicists? They certainty did not represent phlogiston, because phlogiston does not exist. But does that then imply that the models produced were not cases of representation?

On the *relational view* it seems that the phlogiston representations do not satisfy the conditions of representation, because the source (the structure of the model of phlogiston, for example) could not relate to any target (the (observable) phlogiston structure itself).[24] But what about on the *functional view*? It was certainly the case that the models and theories of phlogiston were used in one way or another to demonstrate, interpret, make inferences etc., but does that mean we would be right to conceptualise these models and theories as representations? The answer to this question is unclear on the *functional view*. And this only serves to reinforce the claim that the *functional view* is ill-equipped to set out the sufficient (and perhaps even necessary) conditions of representations.

For many advocates of the *functional view* of representation this vagueness in regards to the necessary and sufficient conditions of representation is a desirable state of affairs. Suarez (2004), for instance, presents his inferential conception of representation as a way for us to:

---

[24] Some advocates of the *relational view* will point out that there could still be some (partial) relation between the source and the target so long as we think of the target in purely structural terms and relinquish the linguistic-device 'phlogiston'. This is the central *structural realist* argument against the *pessimistic meta-induction* (there is no time to discuss this here, but for those who are interested see Ladyman, Ross et al., 2007 and French, 2014 for discussion of structural realism; and see Lauden, 1981 for the canonical presentation of the pessimistic meta-induction).

> [abandon] 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.

Because

> Representation is not the kind of notion that requires, or admits, such conditions. We can at best aim to describe its most general features (Suarez, 2004: 771)

What's more, he argues that:

> On the inferential conception […] there is absolutely no difference in kind between fictional and real-object representation—other than the existence or otherwise of the target.

So Suarez would be happy to conceptualise the phlogiston 'representations' described above as representations, even if the object they purported to represent was merely a fiction.

Elgin too would be happy to call the phlogiston 'representations' actual representations, even after it has turned out that they were inaccurate in some way. This is because Elgin proposes that a

> representation may be accurate or inaccurate. Its claim to objectivity turns not on its accuracy, but on its relation to reasons. (Elgin, 2009: 88).

But now we want to ask: what do we mean by 'reasons' here? And, of course, we find ourselves back with the various cognitive activities of agents including demonstration, interpretations, and inferring.

The central problem of the *functional view* of representation is, therefore, as follows. The *functional view* of representation forces us to give up on the project of specifying in no uncertain terms what does or does not count as a representation. In principle, then, we end up with a situation in which *anything* that can be made use of in the cognitive tasks of agents *could* be conceptualised as a representation. And so we are left with no clear demarcation of what is and is not a representation. It follows from this that we can formulate no worthwhile theory of representation, because we are left only with the possibility of considering the platitudes that the predicate "represent" obeys at the level of practice, disregarding any deeper, or more substantive, account of its nature For some this is completely fine and even desirable (Suarez, 2015). But if this is really where the *functional view* leads, then it does not seem to be an attempt to get a clear '*view*' of representation at all, since it results in the implication that no '*view*' is possible or in fact to be desired. But then why are we focusing on the concept of representation in the first place? And what should we do with all the theories devised in the context of both the *relation* and *functional views* of representation?

Conceptualising Representation as a Structural Relationship

My answer to the final two question at the end of the previous sub-section is that we should develop novel conceptualisations of representation that move beyond the dichotomy between the *relation* and *functional views* of representation.

I am not the first philosopher to posit such an idea. Chakravartty (2009) argues, for instance, that we should overcome the 'dichotomy between' the *relational* and *functional* view of representation by first identifying that the dichotomy itself rests on a 'conflation of means and ends' (Chakravartty, 2009: 13).[25] Chakravartty states that the dichotomy emerges from:

> a conflation of thinking about what scientific representations *are*, as a means to realizing their functions, and thinking about what we *do* with them. [Relational] theories focus primarily on the question 'what are scientific representations?', where representations are conceived as knowledge-bearing entities, such as theories, models, simulations, and diagrams. Functional theories focus primarily on the question 'what is scientific representation?', where 'representation' is conceived as a set of knowledge-exercising practices, constituted by whatever it is that scientists do when engaged in the process of representing things. These are two clearly related, but different questions, and it should come as no surprise that appropriate answers to these questions are clearly related, but different. Naturally, an account of knowledge-bearing entities emphasizes the relations in virtue of which knowledge is borne by those entities. And just as naturally, an account of representational processes emphasizes the various practices in virtue of which that knowledge is exercised. These are complementary questions and

[25] The reader should recall that Chakravartty himself distinguishes between functional and informational views of representation. I have chosen to replace informational with relational so as not to confuse Chakravartty's classification with the informational account of representation I will give in chapter 4 of this essay.

> answers, both contributing to a general understanding of scientific representation. There is no dichotomy between information and function.

So Chakravartty's point is that just as categorical and dispositional descriptions can reveal 'complementary features' of the nature of a given property (Ibid.). Analogously, the *relational* and *functional* views of representation illuminate different aspects of one and the same thing: the nature and function of representation. The *relational* and *functional* view do this, according to Chakravartty, by spelling out on the one hand what properties representations should have to be functional, and on the other hand what functions representations should have as a consequence of their properties.

But the question then arises: Why can we not bring these two perspectives together? In other words: Why cannot we not capture the means and ends associated with representation in one complementary conceptualisation of representation? The problem comes when we try to think about how such a complementary conceptualisation of representation would map onto real concrete instances of representation. Take the phlogiston 'representations' discussed above. The question that seems to call for an answer in this case is: Are the models, theories, laws of phlogiston to be considered genuine representations or not? And it is here that the idea of a complementary conceptualisation of representation seems to flounder, because there appears to be no fact of the matter about how accurate – if at all – a putative representation must be in order to count as a genuine representation.

Those who prioritise the properties of scientific representation – the 'ends' Chakravartty refers to – will say that in order for us to have a genuine case of representation there must be something that the representation is about, where being about something is a relation between two real things (phlogiston and theory, for example). However, those who prioritise the function representations have – the 'means' Chakravartty refers to – may say that in order for us to have a genuine case of representation there must be something that the representation is about, but it need not be the case that what the representation is about is a real thing so long as the representation has a use (in other words, the representation of phlogiston are genuine even if they are about the non-existent entity phlogiston). And this leads Chakravartty to argue that when we try to conceptualise representation we should take in account the fact that what different philosophers mean by "representation" is a 'matter of convention' dependent upon 'the various philosophical uses to which [the concept representation] is put'. (Ibid.: 15).

For my part, I think that Chakravartty's conclusions are based on a fundamental error in conceptualising what representations *are* and *do*. I propose that thinking of representations as on the one hand a thing that is a bundle of properties (knowledge-bearing, relational), and on the other as a thing which permits a bundle of potential functions (inferences, exemplifications), is mistaken from the start. The reason for this mistake, I think, derives from the inclination to think that the distinction between properties and function is one in which representations are *either* taken to be 'in' the world *or* are taken to be 'about' the world.

When, as a matter of convention, representations are taken to be 'in' the world, philosophers look to conceptualise what real, 'worldly' properties all genuine representations must possess. And this leads to a *relational view* of representation in which representations are conceptualised as being identical to one property in particular; i.e. relations. But when, as a matter of convention, representations are taken to be 'about' the world, philosophers look to conceptualise what all genuine representations must be able to do. And here we end up with a *functional view* of representations in which representations are conceptualised as being useful tools in the semantic activities of intentional agents.[26]

Now, I propose that the idea that there is an incompatibility between being 'in' the world and being 'about' the world is a remnant of philosophical dualism; and I reject dualism entirely. For this reason, I propose that the semantic activities of intentional agents are as much a part of the world as tables, chairs, the exhibits and exhibit descriptions at NEMO, the general theory of relativity, or even, perhaps, the number '2' and the concept of beauty. So, on my view, just because a representation is 'about' the world and is used accordingly by intentional agents, does not imply that it is not also 'in' the world. And, *mutatis mutandis*, just because a representation is 'in' the world and possesses certain properties accordingly, does not imply that the representation is not also 'about' the

[26] I think that these considerations at least point to the idea that the debate about how we should conceptualise representation was prefigured by the debate between (scientific) realists and (scientific) anti-realists in many important respects. Spelling out exactly how the two debates relate, however, would require a work of its own, and so I do not intend to explore the idea any further in this essay.

world. Thus, the problem becomes: How are we to conceptualise representation as both 'in' and 'about' the world *at the same time*?

Chakravartty's answer to this last question is that we cannot conceptualise representation as both 'in' and 'about' the world *at the same time*, and so must fall back on different conventions for the use of the word "representation". But I, for one, do not wish to place my trust in the conventional uses of "representation" by contemporary philosophers. And so I offer a novel conceptualisation of representation as follows:

[REP] Representation is the objective structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented (see figure 3 below).

By semantic activities of intentional agents, I mean the set of practices that constitute what agents do when representing things (infer, demonstrate, etc.) By representation-device, I mean the source or vehicle of representation in Suarez (2004) terms (this includes, for example, exhibits at NEMO, paintings in an art gallery, and, perhaps, scientific theories themselves). And by the thing represented, I mean the target of a representation, whether that be phlogiston, electrons, or the Queen of England.

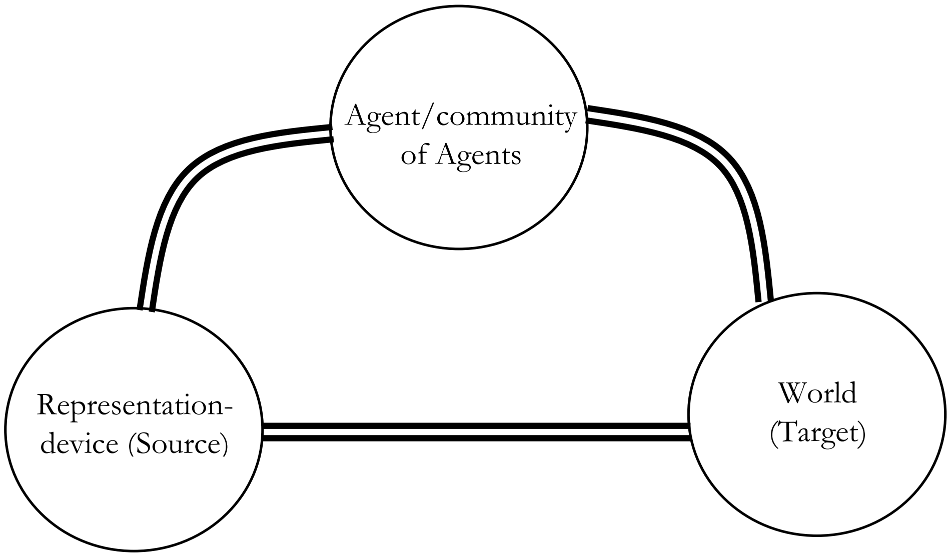

Representation *is* the objective structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented.

Representation =

Figure 3.

The term structural relationship used in [REP] is also of fundamental importance. I use the word 'structural' to express the idea that all three components are interrelated and support one another as one organisationally interconnected whole: there is *one* structural relationship with three components. I use the word 'relationship' to express the idea that we are not just speaking of lifeless relations in the mould of isomorphism, but, rather, we are speaking of a state of relatedness between the agent, world, and representation-device as commonly conceived.[27] Crucially, however, the relationship does not obtain as a consequence of (the semantic activities of) intentional agents any more than it obtains as a consequence of the representation-device or the thing represented. It is a

[27] As we will see later when I talk about *making sense* in my final chapter, it is this agent-world character, I think, that confers onto representations the possibility of being rejected, altered, stored, and transmitted; in short, their defeasible and sharable nature.

relationship, therefore, constituted by agent, world, and representation-device equally. This is the cornerstone of [REP]'s holistic anti-dualism.

The most important feature of [REP] is that it conceptualises representations as both 'in' the world and 'about' the world *at the same time*. Representations, according to [REP], are 'about' the world, because the structural relationship that *is* representation is only possible when intentional agents are able to act in connection with a target (the thing represented). Put differently, conceptualising representation as the structural relationship between agents, the representation-device, and the thing represented presupposes the use of representation by agents. And so [REP] presupposes that, for semantic agents, the sources of representation can function as the basis of inference, demonstration, interpretations etc.

But according to [REP] representations are also 'in' the world to the extent that [REP] conceptualises representations *as* an objective structural relationship between the aforementioned components. Put differently again, [REP] conceptualises representations as something that is enacted as a *real* consequence of the relationship between agent, world, and representation-device. A representation, then, according to [REP], is an objective feature of the world.

So why is [REP] to be preferred to any of the *relational* or *functional* accounts of representation I have considered up to this point? The answer, I think, is clear: because [REP] avoids the shortcomings of both the *relational* and *functional views*. That is, [REP] avoids the central problem with the *function view* by offering a substantive account of the nature of representation as a structural relationship. And [REP] avoids the central problem with the *relational view* (the non-sufficiency argument) by factoring in the importance of *use* or *function* of representation (or at least the sources of representation) in accordance with the semantic activities of agent. I anticipate that the central critique of [REP] will be that it just slices the *relational view* and *function view* together and claims to have resolved the problem. But this is to misapprehend the most important aspect of [REP]: it's anti-dualism. [REP] is an attempt to illustrate that there is no slicing required so long as we view agent and world as part of one system, because agent and world could never actually come apart. Thus, [REP] is an attempt to show that both the *relational* and *functional views* get off on the wrong foot; a foot, that is, that presupposes that we can *either* have an agent-based conceptualisation of representation *or* a word-based conceptualisation of representation, but not both. [REP] is a rejection of this dichotomy out the outset. As such, [REP] is motivated by the problems inherent to the *relational* and *functional views* and is an attempt to dissolve these problems by arguing that they emerge from dualistic presuppositions. According to [REP], these problems simply do not arise when we undermine dualism and so [REP] answers to the problems of the *relational* and *functional views* by demonstrating that they are pseudo problems, concocted by philosophers struggling to free themselves from the shadow of Descartes.

[REP], however, does have one serious problem of its own, because the question of whether or not the thing represented also has to *really exist* in order for the structural relationship to be objective is left open. And without an answer to this question, critics of [REP] can simply argue that [REP] only implies that representations are 'in' the world when there is actually some *thing* (target) that a source relates to, because otherwise we do not have all three necessary components between which representation – as a structural relationship – obtains. Hence, [REP] is just another species of the *relational view* discussed above. If we return to the example of the phlogiston 'representation' above, therefore, critics of [REP] will simply argue that [REP] must conclude that the models, theories, and laws of phlogiston are not genuine representations, because there was no *thing* to be represented (and hence it was impossible for there to have been a structural relationship between all three components necessary for [REP]). And so [REP] merely falls onto one side of the conventional divide proposed by Chakravartty (2009): the side that uses "representation" to mean a relation between actually existing things.

My counter argument to this accusation will take the remainder of this essay to elaborate.[28] But, suffice to say, I think this critique of [REP] trades on a flawed conceptualisation of what [REP] means by a structural relationship; that is, a flawed conception of the nature of the relationship between world, agent, and artefact (representation-device). To anticipate, I will argue that critics of [REP] must come to realise that the question of the reality or unreality of the thing represented (the target) that [REP] speaks of is has no impact on the question of the reality or unreality of the structural relationship between agent(s), world, and aretfact. In this way, even if we can be unsure about the nature of the components that sit in the structural relationship this *is* representation, the structural relationship itself remains real and objective. In short, the structural relationship formulated in [REP] exists independently from questions about the reality or non-reality of its components, because the structural relationship *is* the

[28] My final response will come in my conclusion below.

dynamic interaction of its various components; it *is* the dynamic and creative interaction between agent(s), world, and artefact(s).

On my view, therefore, representation *is* a structural relationship between agent(s), world, and artefact(s). And what the various components of the structural relationship *do* is bring representation into being. Representation, then, is the *enacted* consequence of a relationship between the three aforementioned components. In this way, my [REP] conceptualisation of representation should be embedded within the *enactivist* movement in philosophy. Enactivism is the philosophical position that holds that:

> Organisms do not passively receive information from their environments, which they then translate into internal representations. Natural cognitive systems […] participate in the generation of meaning […] engaging in transformational […] interactions: *they enact a world.* (Di Paolo, Rhohde, De Jaegher, 2014: 33)[29]

It should be noted, however, that [REP] does not assent to the claim that organisms enact anything when taken alone, for this would once again open the door once again to a form of dualism. Instead, [REP] holds that representation is enacted as the result of a relationship between organisms, the world, and artefacts that issues from the agent, the world, and the artefacts. So, according to [REP], representation is the enacted result of a dynamic interaction between an acting organism, its environment, and

[29] For more detailed accounts of enactivism see, for example, Clark and Toribio (1994) and Thompson (2007).

relevant artefact, which privileges neither the organism, the world, nor the artefact.[30]

For now this description of representation as an enacted structural relationship will have to remain vague and, I am sure, unsatisfying. I ask the reader to postpone judgement until chapter 4, where I take a novel perspective on the structural relationship in question. Firstly, however, I must consider what [REP] implies for one very particular species of representation: the representations of science.

[30] I must stress at this point that my use of enactivism is an aberration from the conventional usage in the enactivist tradition whereby the focus is on a particular view of (the emergence of-) cognition. I do think, however, that my use of enact, enacting, enactivism etc. throughout this essay is consistent with the basic enactivist anti-dualist idea that some *x*, say, mind or representation, is generated or brought into being as the result of an interaction between organism and world. The only difference, then, I propose, is that when I make use of the term enact I am thinking about a second order enacting of *representation* that is the result of an interaction between cognitive organisms, the world, and artefacts, and not a first order enacting of *cognition* that is the result of an interaction between organisms and world. For those who find this representational extension of enactivism intolerable, I ask simply that for enact(s), enacting, etc. they read generate, induce, bring about, or some other such synonym. My overarching argument will not be affected by the substitution of such synonyms for enact, enacting etc.

Chapter 2.

Scientific Representations

In this chapter, I consider *scientific* representations. Firstly, I distinguish scientific representations from other kinds of representation by reference to the criteria of giving the nearest approximation of reality possible.. Then, I propose that all scientific representations have models as their representation-device (or source). This leads me to a consideration of the way that models are constructed and used. And, subsequently, to the assertion that the content of models cannot be understood without reference to the way they are constructed and used. Following from this assertion, I argue that the content of models must not be taken to be 'in' the representation-device alone, but, rather, should be understood as being in the logical space that the representation-device present. Finally, I argue that the content of models is the presented result of the structural relationship between agent, world, and representation-device that aims to give the nearest approximation of reality possible. I argue that this special structural relationship *is* scientific representation.

What's Special about Scientific Representations?

We find representations in a great many domains of human activity. In art, for instance, we find representations of people in the form of paintings, of concepts in the form of sculptures, of fictional stories in the form of ballet performances, and of emotions in the form of poetry. Just as in a the domain of a child's play we find representations of forts in the form of collections of cardboard boxes, of aeroplanes in the form of four poster beds, and of bears in the form of carefully sewn material stuffed with cotton. But the representations of one domain in particular will occupy my considerations in this chapter: the representations of science.

It would be prudent, then, to start by giving a definition of what I mean by *science*. I follow Ladyman, Ross, et al. (2007) by conceptualising science in a pragmatic manner that refuses to postulate a 'particular set of positive rules for reasoning that all and only scientists do or should follow' (Ladyman, Ross, et al., 2007: 28). In this way, I also endorse their claim that 'there is no such thing as a 'scientific method''. It follows from this, as Ladyman, Ross, et al. (Ibid.) point out, that science is:

> demarcated from non-science solely by institutional norms: requirements for rigorous peer review before claims may be deposited in 'serious' registers of scientific belief, requirements governing representational rigour with respect to both theoretical claims and accounts of observations and experiments, and so on.

The institutions that Ladyman, Ross, et al. are referring here are said to be the 'institutions of modern science'. For their part, Ladyman, Ross, et al. are content to leave what is meant by an institutions of modern science 'open to the rational judgements of observers of institutional processes' (Ibid.). They do stipulate, however, that scientific institutions in their sense must be 'anchored around attempts to determine the objective structures in nature', as opposed to other 'anthropocentric' or 'ideological' ambitions.

Taking this into account, the definition of science endorsed by Ladyman, Ross, et al. becomes the following:

> [SCI] Science just *is* our set of institutional error filters for the job of discovering the objective character of the world—that and no more but also that *and no less* (Ibid.).[31]

I endorse [SCI] as a conceptualisation of science. My concern now, however, is not so much with what science *is*, but rather with what science *does*. Back in the introduction, I stated as my second assumption that one major aim of science is to produce representations of the world. But now, armed with the conceptualisation of representation I developed in chapter 1, I must consider in more detail what this assumption amounts to.

Given [SCI] and my second assumption, I will call a scientific representation any representation that has been produced in accord with

---

[31] I assume that some may balk at the idea that science is conceptualised as a collection of normative institutional error filter. And Ladyman, Ross, et al. are alive to this criticism. They argue, however, on pragmatic grounds again, that science has proven that it allows for us to 'achieve significant epistemological feats by collaborating and by creating strong institutional filters on errors'. They do accept that 'a person who finds science unimpressive, or demeaning (think of Coleridge or Wordsworth), or Faustian, will have no reason at all to be persuaded by [their arguments]'. And this would also hold for any critique of the institutions of science in a Marxist or Foucauldian fashion. It will not be my purpose to enter into these kinds of debates, however. I will, therefore, adopt [SCI] as a matter of heuristics. Thus, the question of what constitutes a *truly scientific* institution is not one I will consider here. That is not to say that this question is entirely disconnected from my project, but only that any answer one gives to this question one way or the other will have no bearing on my eventual arguments for how we should conceptualise the nature and function of scientific representation.

the appropriate scientific institutional norms and error filters. But this, of course, does not tell us much about scientific representations unless we can say what the purpose of such institutional norms and error filters is; that is, unless we can say why we have such institutional norms and error filters in the first place.

The answer Ladyman, Ross, et al. give is clear: the institutional norms and error filters have developed in such a way because they provide 'feeble' epistemic agents like us the best chance of forming the most defensible beliefs about reality (Ibid.). As they say:

> we assume that the institutions of modern science are more reliable epistemic filters than are any criteria that could be identified by philosophical analysis and written down. Note that we do not derive this belief from any wider belief about the reliability of evolved human institutions in general. Most of those—governments, political parties, churches, firms, NGOs, ethnic associations, families, etc.—are hardly epistemically reliable at all. Our grounding assumption is that the specific institutional processes of science have inductively established peculiar epistemic reliability (Ladyman, Ross, et al., 2007: 37).

But what, then, is this assumed peculiar epistemic reliability that justifies why we have scientific institutional norms and error filters in the first place? The answer, I think, can be put like this: scientific institutional norms and error filters are assumed to lead to a peculiar epistemic reliability because they are assumed to facilitate the best approximation of reality possible. It is for this reason that Ladyman, Ross, et al. argue that 'With respect to anything that is a putative fact about the world, scientific institutional processes are absolutely and exclusively authoritative' (Ibid.: 28).

As a consequence of [SCI], then, to say that the aim of science is to produce representations of the world is just to say that the aim of science is to produce representations that give the best approximation of reality possible. Scientific representations, therefore, are supposed to be the most accurate approximations of reality possible. So scientific representations are supposed to be our best attempt to say how the world really is. In this way, scientific representations are ordinarily differentiated from representations in the domains such as art or child's play. Thus, by accepting [SCI] I am committing myself to the claim that scientific representations are special, because they aim at giving the most approximately accurate account of reality available.

Models and Selective Resemblance

In chapter 1, I advanced my own conceptualisation of representations as the structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented ([REP]). [REP] is supposed to be a conceptualisation of all kinds of representations, whether in the domain of art, science, or child's play. But if this is the case, then it would seem we require some justification for the [SCI]-based claim that science produces representations of a qualitatively different kind from representations in other domains. That is, some justification for the assumption that scientific representations, unlike other kinds of representation, are a particular form of structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented *aimed at giving the best approximation of reality possible*.

Nelson Goodman famously dismissed the idea of marking any such qualitative difference on the grounds 'that in science, unlike art, the ultimate test [of a representation] is truth' (Goodman, 1976: 262). In fact, Goodman contends that 'truth by itself matters very little in science' and that ultimately we should view the truth of representations in all domains in terms of a 'matter of fit' between the representation and the world (Ibid.). This may be right if we try to mark the qualitative difference between scientific and other kinds of representations in term of the *actual* truth-value of scientific representations.[32] But we could instead attempt to mark the difference in terms of the *aim* of scientific representations; i.e. the reliability of the truth-value of scientific representations. I will defend this second alternative. To do this, I will now examine the construction and use of scientific representation-devices: scientific models. I will do this because an examination of the way models are constructed and used gives the clearest indication of the aim of representational practices in the sciences.

To begin we should first note that all scientific representations make use of models; whether this is a model of an historical event, the carbon cycle, or the structure of the atom. Philosophers of different persuasions argue over whether we should conceptualise models as physical objects, such as the foam model exhibit of the human cell at NEMO (Schaffner, 1969; Morgan and Boumans, 2004); as abstract

[32] I do necessarily mean to commit myself to a Goodmanian claim about the relation between representations and truth, but mean only to leave the question open.

objects, such as a set-theoretical structure or a mathematical equations (Suppes, 1960, 1989; da Costa and French, 2000); or as fictional objects, such as a frictionless pendulums (Barberousse and Ludwig, 2009; Frigg, 2010a, 2010b; Godfrey-Smith, 2006, 2009; Leng, 2010; Toon, 2010).[33] I have already said that scientific representations *are*, at basis, structural relationships of the type given by [REP]. As such, I hold that models should only be conceived of as one component of the structural relationship that *is* representation: the representation-device (the source or vehicle of representation). On my view, therefore, the intensity of the debate about whether models should be conceptualised as physical, abstract, or fictional objects is predicated on a mistaken assumption about what such a debate reveals. The mistake, that is, of assuming that such a debate will help us to determine what representations *are*. I hold, in accordance with [REP], that all such a debate will merely help us to determine is what representation-devices, i.e. models, *are*.

The ontology of representation-devices will not concern me any longer here therefore. But I am interested in thinking about what representation-devices do (from here on I will say "model" in place of "representation-device", but the reader is asked to keep in mind that I am not therefore equating the term "model" with scientific representations as conceptualised by [REP]). One thing that models do is quite clear: they model phenomena. An example of a model performing this function

[33] Where any one philosopher draws the line between differing accounts of what models *are* is invariably dependent upon their metaphysical commitments. For example, a nominalist may put abstract models in the same class as fictional models, whereas a realist/Platonist will take abstract models to be a separate class from fictional models (and may even reject fictional models all together (c.f. Giere, 2009 and Pincock, 2012)).

would be the model of Rotterdam's Erasmus Bridge in NEMO science museum. But it is also be the case that all of the other models in NEMO – that is, the models of human cells, DNA molecules, and space-time itself – are also models of some phenomenon or at least putative phenomenon.[34] Still, there appears to be considerable differences between the metal model of Erasmus Bridge, the foam model of the human cell, and the mathematical model of the potential topology of space-time. And so we feel drawn towards the question: How, exactly, do models as diverse as the ones found in NEMO model phenomena?[35]

I will follow Van Fraassen (2008) by positing that models like the ones in NEMO model phenomena 'By selective resemblance and selective (even systematic) non-resemblance' (Van Fraassen, 2008: 35). What is meant by resemblance here 'need not consist in sameness of properties, but can also be at higher levels' (Ibid.). So:

> Resemblance may consist in having properties in common, or instead in having properties that have properties in common with relevant properties in what is represented. That is, the representor may have properties which form a structure resembling a structure formed by the properties of the represented, and so on up in a hierarchy of types (Ibid.:34)

[34] This claims assumes a substantivalist conceptualisation or spacetime as opposed to a relationalist conceptualisation, and so is, potentially, open to critique (such a critique would also, I assume, be offered by Neo-Kantians who take space and time to be features of our understanding). For my purposes, however, these debates are irrelevant and should be settled by the mechanisms of science ([SCI]).

[35] It is imperative that the reader does not confuse this with the question: What *is* a model? For this would lead us into the discussion about the ontology of representation-devices, which I have already left behind.

Now, the point that Van Fraassen is at pains to make explicit is that resemblance is selective. By this he means to say that a model only aims to resemble some features or aspects of the target phenomena.

Take, for example, a *scale* model like the Erasmus Bridge model in NEMO. Van Fraassen notes that in the cases of such models 'scaling is not simple reduction or increase in size in all dimensions', because 'useful scaling trades not just on the obvious resemblances in shape but on distortion, both resemblance and non-resemblance being selective in a way dictated by the purpose at hand' (Ibid.: 49). To consider what Van Fraassen means here consider the following example. Say we want to compare the strength of a cylindrical beam with constant density that is part of the actual Erasmus Bridge in Rotterdam with the strength of a cylindrical beam with constant density that is part of the model of Erasmus Bridge in NEMO. The strength – both compressive and tensile – of a cylinder is proportional to its cross-sectional area. But the mass of a cylinder is proportional to its volume. Thus, the strength to mass ration of the cylindrical beam that forms part of the Erasmus Bridge model in NEMO will not correspond to the strength to mass ratio of the cylindrical beam that is part of the actual Erasmus Bridge in Rotterdam, *when the scale model has been [constructed] by a process of geometric or dimensional scaling.*[36] So not *all* of the properties of the target phenomena (the actual Erasmus Bridge) are preserved by a *geometric* scale model (the model in NEMO).

[36] These observations are supported by the square-cube law as formulated by Galileo, which implies that as an object is isometrically scaled up or down its volume increases or decreases at a greater rate than its surface area.

The point, then, is that a model like the Erasmus Bridge model in NEMO must be selective about which properties to resemble if it is to function as a helpful model at all. We can appreciate the necessity of this 'inevitable occlusion or distortion' involved with modelling if we imagine a scaled *up* model of the Erasmus Bridge; i.e. a model of the Erasmus Bridge that is much bigger than the actual Erasmus Bridge in Rotterdam. In this imaginary scenario it is obvious that a scaled up model of the Erasmus Bridge would have to selectively resemble the actual Erasmus Bridge, because given a perfect resemblance of *all* the properties of the Erasmus Bridge the scaled up model would have a strength to mass ratio that would cause the cylindrical beams of the model to break down under their own weight.

This example of the Erasmus Bridge model is designed to show that a scale model, to be a functioning model at all, must *selectively* resemble the properties of its target. And as the example of the Erasmus Bridge model indicates, this selective resemblance can only be achieved when a model shares some properties with its target, but distorts others. But we noted earlier that the Erasmus Bridge model seems, *prima facia* at least, to be a very different kind of model than the painted mathematical model of Einstein's Field equations also housed at NEMO. However, Van Fraassen also argues that 'mathematical [modelling] too involves in general necessary or inevitable distortion, in both simple and subtle ways' (Ibid.: 40).

The central claim that Van Fraassen makes is that mathematical models all involve simplifications, of one kind or another. In yet another illuminating example, he argues:

> Consider this beer glass on the table: each has a *shape*. What this shape is, precisely, we do not know. In the heady early days of the mathematization of our world picture it could be assumed that this shape is [described by] a precise function of the spatial coordinates. The edge of this table could be thought of as a straight line, hence [described by] a function of form $y = ax + b$. But of course, the edge of a table is not perfectly straight…If eventually the table, the beer glass, and their environment are re-conceived as assemblies of classical particles, they still occupy precise regions of space. These regions may be similarly [modelled] by functions on the spatial coordinates. It may be a bit arbitrary exactly which particle assembly is the glass at any given time – but upon any such arbitrary, admissible choice, the table and glass have a definite shape. So now these objects can be [modelled] by a mathematical model *in the same way but more accurately than before* […]
>
> We are speaking here of the continuum of classical mechanics which has equal use for the [modelling] of each primary quality: length, duration, shape, size, number, mass, velocity, what have you. The equation of the primary quality *shape* with *geometric shape* is in effect the assertion that a certain [model] is completely adequate. But now we must ask: what exactly is this [model]? Not only the question as to what shape the glass has, but *that* question is continually answered differently. (Ibid.: 46)

The point Van Fraassen is making here is that when we come to evaluate the credentials of a given mathematical model, we are inclined to prefer the mathematical model that is consistent with our own scientific image, because 'in retrospect [we] look upon the scientific image inherited from the older generation as *open*, *vague*, *ambiguous* in the light of our new understanding (that is, in the light of alternatives not previously conceived)' (Ibid.: 47). In other words, we come to see that the simplifications that buttressed the scientific image of the 'older generation' are unsustainable in the light of our own scientific image (consider how the simplification of the edge of the table being straight is able to be removed when we re-conceive of the table as an assembly of particles occupying a region of space). And thus we suppose that mathematical models previously deemed adequate are less accurate than our current best mathematical models.

But now we must ask: Why should we prefer our scientific image over its predecessors? And why should we suppose that our current best mathematical models gives a truer story of the world than their predecessors? Van Fraassen's aim is to make us see that providing an answer to these questions is tricky, because *all* mathematical models involve some kind of 'distortion', resulting from the way the properties of the target system are modelled (Ibid.: 40). In one model this 'distortion' may emerge from the simplification that the shape of a beer glass can be given by *supposing* that the table edge is straight. In another model the 'distortion' may emerge from the assumption that the shape of the beer glass is a particular assembly of particles in a region of space at time t. In *both* of these models, however, 'distortion' is at play: the table may not be straight and the 'shape' of the beer glass at time t+1 *may* not accord with the function that describes the assembly of particles at time t. And so a model's resemblance of the shape of a beer glass must be *selective* in both instances: either we get a function that gives the shape of a beer glass given the simplification that the table edge is straight, or we get a function that gives the shape of a beer glass given an arbitrary, but admissible, assumption about which assembly of particles counts as the glass at any given time *t*.

Some may think that the example of the shape of a beer glass is too specific to substantiate the claim that all mathematical models involve selective resemblance. They may argue that it is the concept of *shape* that lacks clarity in the example under consideration. But simplifications like those discussed above proliferate in nearly all mathematical models constructed in the mathematical sciences. In astrodynamics, planets are standardly taken to have only mass and shape, ignoring all other properties (Cartwright, 1989; Ch. 5). In neo-classical economics, markets are taken to be in perfect equilibrium and agents are taken to be perfectly rational and omniscient. And in physics and engineering, there is continued talk of frictionless planes, infinite velocities/densities, and point masses (Frigg and Hartmann, 2012). So should we then say that like the concept of shape, the concepts of planet, market, plane, velocity, density, and mass are all equally unclear? This would surely be too much given the centrality of some of these concepts to the sciences in question. Thus, I concur with Van Fraassen's claim that if there is one thing that mathematical models do it is distort the way the world really *is*, because mathematical models only selectively resemble their target systems.

We have now seen that both physical (scale) and abstract (mathematical) models, model phenomena by means of selective resemblance and selective non-resemblance.[37] But models are not exclusively used to model phenomena, they are also used to model data. Data models are also representation-devices in my sense, because they are just one component of the structural relationship that I have conceptualised as representation ([REP]). But they can be differentiated from the models of phenomena I have been discussing in virtue of the fact that they are not models of such things as bridges or beer glasses, but are instead models of raw data. Raw data – more often than not given in numerical form – is the result of immediate observation. Despite this immediacy, however, I still hold that data-models too trade on selective resemblance.

A data model is a standardised, regimented, and often corrected or rectified version of the data that we gain from immediate observation (Frigg and Hartmann, 2012). Familiar data models include: the Body mass index, as derived from data concerning an individual's weight and height; a graph depicting the appreciation or depreciation of the value of the stock of a particular company over the course of a month, constructed from data gathered at a stock exchange at various times throughout the month in question; and the average rating of a new movie on an internet website, constructed from data gathered from various reviews in newspapers, blog posts, and film-studies journals.[38] Data models, therefore, come in the

[37] There is little point in discussing whether or not fictional models would model phenomena based upon selective resemblance, because such models evidently select which properties to resemble and which to not resemble – that's whole point of their being fictional!

[38] For a clear exposition of how data structures are extracted from phenomena by measurement see Muller (2011: especially the schema on page 98). This will also be particularly pertinent to the models, measurement, and logical space sub-section later in this chapter.

form of, i.e., graphs, distributions, or numerical values (c.f. Mayo, 1996; Harris, 2003).

What's important for my argument here, however, is how the data gathered is organised into a data *model*. This organisation – which often takes the form of a summarising, standardising, or regimenting – is crucial to understanding what a data model does. The reason for this is because the data model, far from being an unrefined presentation of the raw data, is a 'processed artefact' (Van Fraassen, 2008: 167). What this means is that the data model is constructed from the raw data in such a way 'so as to replace the relative frequency counts [of the raw data] by measures with a continuous range of values' (Ibid.). In other words, a data model combines and transforms raw data in the process of giving an analogue (continuous), as opposed to digital (discrete), presentation of that raw data. As Suppes (1962: 253) put it:

> It is common for models of a theory to contain continuous functions or infinite sequences although the confirming data are highly discrete and finitistic in character.

And so I argue that the frequencies of raw data that we find in nature by means of observation are modelled in a distorted manner.[39]

So do the distortions of data models count as simplifications or as selective resemblances? I think that the answer is clearly: yes, they count

[39] It is worth noting that Van Fraassen (2008: 167-168) makes a difference between 'data models' as summaries of frequencies of raw data, and 'surface models' which exemplify the distortions I have been referring to here. He also says, however, that he 'will not insist everywhere pedantically on this distinction', and for my purposes it is not necessary to introduce further complications at this stage.

as both simplifications and selective resemblances. The reason, again, is because data models typically strive for averages, for 'smooth' curves through graphs, and so for the exclusion of outliers and anomalies (Van Fraassen, 2008: 167). The data model, then, is an artefact constructed by a process of organisation that values 'neatness' and standardisation above all else (Suppes, 1962). And this requires a simplification in the way the data is presented, which can only be achieved by presenting the data selectively via the application of sophisticated statistical techniques. In his (2006), Giere made exactly this point when examining how 'astronomical data, gathered in the form of black and white photos, are processed to yield images in color of a nebula' (Van Fraassen, 2008: 168). Giere highlighted that,

> The images presented…are conclusions. These images present a picture that is continuous, or at least fine-grained. The actual data cannot be that fine grained. The data are made up of individual events recorded in various detectors at different times and processed by various physical and computational means. The images are constructed using those data, but go beyond the data. (Giere, 2006: 48)

All this is to say that data models model data in exactly the same way that physical (scale), abstract (mathematical), and fictional models model phenomena: by means of selective resemblance.

What, then, can we take from this analysis of what models do? First and foremost, we can appreciate the fact that models – as the representation-devices of scientific representations – do *not* aim to model the world as it is apprehended by us in experience. Nor do they aim to model the world by being exact replicas of the world (recall my Goodmanian-inspired critique about what representations are not in chapter 1). They do not aim to do these things for a number of reasons. All of these reasons boil down to one factor: utility. Data models, for example, do not typically aim to model data as it is observed and collated, because a model of this sort would be unwieldy, inaccessible, and ultimately unhelpful to the theorist. Similarly, models of phenomena do not aim to be model replicas or copies, because such models would fail to be of use in describing or presenting the phenomena in question (consider the Erasmus bridge example above). The act of modelling then is indexed, in both its data and phenomenal varieties, to a purpose of utility. And it is this purpose of utility that guides the selection of salient properties to be modelled.[40] Thus, as Van Fraassen suggests, we should come to see 'The "selective" in "selective resemblance" [as] a delicate, highly nuanced, contextually sensitive qualification' (Van Fraassen, 2008: 57).

In short, then, models are representation-devices that selectively resemble the properties of a target system (bridges, raw data) for a purpose of utility: models are processed artefacts that function as a device that can be both used by agents and that at least selectively capture features of the world. The selective resemblance may obtain at the level of individual properties (the strength of cylindrical beams in the Erasmus Bridge) or at the level of structure (smoothing of raw data in a graph). But the question then becomes: What are the constructive processes by which models are made to selectively resemble their target systems? To answer this question,

---

[40] The phrase, purpose of utility, may seem to be a pleonasm, but I will continue to make use of it to stress that the utility of a model – as a measure of preferences – is related to the purpose of a model – as fixed by the agent.

I will now briefly relate what I have been saying to the discussion of *abstraction* and *idealisation* in contemporary philosophy of science.

Abstraction and Idealisation

Abstraction and idealisation are two processes by which models are constructed.[41] That is, in my terms, two processes by which the representation-devices of scientific representations are constructed.

> Abstraction, according to Chakravartty (2007: 190):
>
> is a process in which only some of the potential many relevant factors present in reality are represented in a model […] concerned with some aspects of the world, such as the nature or behaviour of a specific object or process. Here one excludes other factors that are potentially relevant to the phenomena under consideration.

Abstraction, then, is the process of 'stripping away' the properties of a target system that are not relevant to the modelling ambition at hand; i.e. abstraction is a selection (Cartwright, 1989: Ch. 5). By doing this a model is able to focus on as limited a number of relevant properties as possible. This is helpful in one important respect: models are constructed that best serve our explanatory purposes. This is the case because abstraction facilitates the construction of refined models that do not attempt to take into account all of the potentially relevant properties of a target system. In this way, factors that are deemed to be of negligible importance are

[41] Some philosophers call abstraction, 'Aristotelian Idealisation', and idealisation, 'Galilean Idealisation'. (c.f. Musgrave, 1981).

excluded from the model; i.e. are not modelled. And thus the constructed model itself is able to give a simplified description of a complicated target system that may otherwise be simply intractable.

A good example of abstraction is the model of planetary motion I briefly mentioned above. The model in this instance describes the motion of the planets by considering only two of their properties: shape and mass. Planets, however, have a number of other properties. Axis tilt, for instance, is the tilt of the planet's rotation axis with respect to its orbital plane. And inclination is the tilt of a planet's orbit with respect to its ecliptic. These properties – and others not mentioned – do seem to be relevant to our understanding of what a planet *is*, but they are excluded from many models of planetary motion. Here, then, we have a clear case of abstraction, whereby a model excludes certain properties of its target system in a bid to best serve our explanatory purposes. In the case in question, the model excludes the properties of, for instance, axis tilt and inclination, in order that we may best be able to explain how planets move through space and time

Turning our attention to idealisation, we can once again refer to Chakravartty's (2207: 191) characterisation. He says:

> The trademark feature of idealization is that model elements are assembled in such a way as to differ from the things they [model], not merely by excluding factors as in the case of abstraction, but by incorporating factors that cannot exist as [modelled] given the actual properties and relations involved.

Idealisation, then, is the process of deliberately constructing models that include fictionalised aspects. The reason for doing this is once again to allow for models to serve our best explanatory purposes. To allow, in other words, that a model describes a target system – suitably fictionalised in some respects – in a way that is most helpful to our explanatory ambitions.

Examples of idealisation abound in the sciences. In physics, for example, models are constructed of point masses moving on frictionless planes. Biologists construct models describing isolated populations. And economists, enamoured with the idea of *homo economicus*, construct models in which agents are omniscient and markets exemplify rational equilibrium. All of these properties are projected onto the target system by the idealised model; they are not to be found in the target system itself. Instead, they are deliberate fictions designed to facilitate the construction of models that describe the target system in the most useful way possible given a particular explanatory purpose.

We can now link this discussion of abstraction and idealisation back to be discussion of models and selective resemblance in the previous sub-section. We can do this by noticing that both abstraction and idealisation are *processes* by which models simplify or distort their target systems. In the case of abstraction, this simplification or distortion takes the form of an exclusion of potentially relevant properties of the target system. And in the case of idealisation, this simplification or distortion takes the form of our fictionalising the properties of the target system. Both abstraction and idealisation, then, are constructive processes by which models selectively resemble their target systems: either by choosing what properties of the target system to model, or by choosing what some of the properties of the target system should be.

We are therefore brought back to the topic we encountered at the end of the last sub-section; namely, the purposes that guide modelling activity in the sciences. I proposed earlier that modelling is indexed, in both its data and phenomenal varieties, to a purpose of utility. And I think the focus on the modelling processes of abstraction and idealisation serve only to reinforce this claim. The reason for this is clear: because both abstraction and idealisation are processes that guide the construction of models towards our best explanatory purposes. Explanation then, I think, is the most integral utility function of models, and hence it is the desire for more accessible, simpler, and comprehensive explanation that underpins the process of model construction in the sciences.[42] The purpose of utility that guides the selective resemblance of models to their target system can, therefore, be understood as a purpose of explanatory utility through the lens of my analysis of abstraction and idealisation.

But what then, we may ask, is explanatory utility? And how can we compare different models in terms of their explanatory utility? To move towards an answer to these questions, we must shift our attention away from the question of what models do, to the question of how models are used.

[42] It is also the case that models are constructed in the sciences for anthropocentric purposes. In engineering, for example, models are constructed with the purpose of enabling us to build bridges, develop advanced modes of transportation, and produce renewable energy technologies. Still, I think that even in these cases explanation is the motivating purpose of model construction. It is just that the explanation that is deemed useful or helpful in these cases should be explanation that directly benefits, i.e., human civilisation.

Models, Measurement, and Logical Space

I have now made two claims in this chapter. Firstly, that models – understood as representation-devices – model their target systems by means of selective resemblance, which is guided by a purpose of utility. And, secondly, that the purpose of utility that guides the construction of models is primarily a purpose of explanatory utility.[43] In this section, I will consider the use of scientific models. To do this, I will ask: How are models used? And: What makes a model useful or helpful to the model user?

Let's begin with the question: How are models used? My argument here is that the central use of models is to explain. I define a model-based explanation as follows:

[MEXP] A model-based explanation is any description that uses a model to describe some or all of the logically possible configurations of the properties or relations of the target system of a model.

Model-based explanations can, therefore, be mechanistic, causal, nomological, pragmatic, or even unificationist, and so can be subsumed within practically any account of explanation. But this inclusivity of model-based explanations prompts the question: How can we determine whether or not a model is explanatorily helpful or useful to the model user? The answer, I think, is connected to the fact that model construction always follows a measurement.

[43] This second claim was buttressed by my discussion of abstraction and idealisation.

So what does it mean to say model construction always comes after a measurement? One may ask, for example, if a historical model of the battlefield developments in the battle of Waterloo is constructed after a measurement of anything. Just as one may ask what measurement preceded the construction of a pictorial model of the carbon cycle. As these examples illustrate, then, one may intuitively feel wary about supposing that all models are built upon a measurement, because there appear to be a range of models whose construction does not seem to have followed from a measurement at all.

The problem here, I think, is with the ordinary conceptualisation of measuring as the act or process of assigning numbers to a phenomenon. I want to say that although the assigning of numbers to a phenomenon is something that can follow from measurement, the assigning of numbers fails to account for what measurement really is. I once again concur with Van Fraassen who posits that:

[MEAS] Measurement is an operation that locates an item (already classified as in the domain of a given theory) in a logical space (provided by the theory to [model] a range of possible states or characteristics of such items). (Van Fraassen, 2008: 164)

whereby

> The HSB color space, with dimensions hue, brightness, and saturation is a good example of a logical space, but so is the PVT space in elementary gas theory, phase space in classical mechanics, Hilbert space in quantum mechanics; [and] space and time themselves also serve as examples [of logical spaces]. (Ibid.)

Logical space, therefore, may be viewed as a space of possible states of affairs in the early Wittgensteinian sense; i.e. a 'field of possibilities' onto which items can be located (c.f. Wittgenstein, 1922: TLP 1.13, 2.013, 2.0131, and 2.202). Such possibilities could include, for instance, the location of unit on a battlefield in present-day Belgium, in 1815; or the exchange of something between the hydrosphere and the atmosphere on the planet earth. And such items could include be the Scots Greys cavalry unit or the element carbon[44]

Now, if measurement is the location of items in logical space, then it is clear how all models follow from measurement. This is the case because models are constructed to model a target system by selectively resembling the target system in question. And model construction will simply not be able to get off the ground without there first having been a location of the properties of the target system that are to be selectively resembled in a logical space of possible states of affairs. For example, in order to model the solar system it is first necessary to locate the planets in the logical space of celestial bodies, and thereby confer onto the planetary 'items' the possibility of entering into states of affairs permissible within the logical space of celestial bodies. Of course, different models of the same target system can be constructed and these models may follow from the location of the same properties of a target system in different logical spaces – as was the case in the shift from Ptolemy's geocentric model of

[44] There is a very important question regarding what it is that constitutes the 'logical space' in which the item can be located. In the citation I have given, Van Fraassen indicates that he takes it to be the background theory against which the model is constructed which constitutes the logical space. I do not share this view, but I do not want to explain why until I move on to discuss the content of models in the penultimate sub-section of the present chapter.

the solar system to Copernicus' heliocentric model of the solar system. But still we must be aware that the very possibility of modelling a target system at all requires first that the relevant properties of the target system in question have been located in a logical space of some kind. To be clear, this is not to say that the logical space in which some of the properties of a target system are located need always be as specific as, say, the logical space of Hilbert space in quantum mechanics. But in order for a model to be a selective resemblance of something in the first place, it must be the case that some aspect of its target system *have been* located in a logical space of some form. Selective resemblance – and hence model construction – presupposes the location of some of the properties of a target system in logical space. And, therefore, in order to model by selective resemblance, some properties of a model's target system must have been measured according to [MEAS].

This exposition, then, asserts that the very possibility of model construction depends upon a measurement. It asserts that models as representation-devices which are processed artefacts proceed from measurement ([MEAS]). I will say that modelling *begins* from measurement to capture this idea. But what, then, does this imply about the use of models to give model-based explanations ([MEXP])?

Given the history of ideas, we can assume that the first models constructed proceeded from the location of items in logical spaces as general as, say, the logical space of space and time, the logical space of earthly elements, or the logical space of the supernatural. Religious models and, perhaps, the Aristotelian model of physical reality would be examples of models that *began* from the location of items in such general logical spaces.[45] Now, these general logical spaces make permissible a great many possible states of affairs; from omnipotent deities to special weightless and incorruptible elements. And for this reason any model-based explanation ([MEXP]) that is supported by a model that *begins* from the location of items in these general logical spaces will find it very difficult to describe even a fraction of the logically possible configurations of the properties or relations of the model's target system. Why? Because the model in question is a processes artefact that has been constructed following from a measurement that locates the properties of the target system in a logical space that permits of a great many logical possibilities; and so the model itself will be a selective resemblance of the target system's properties as located in a logical space that admit of a huge range of possible states or characteristics of such properties.[46]

My argument, then, is that any model that *begins* from the location of items in such general logical spaces will have a low explanatory utility, because the number of possible states of affairs that any model must try

---

[45] What I am saying here is consistent with the ideas expressed by Wilfrid Sellars (1962) regarding the difference between the manifest and scientific images of man. I will not explore this line of thought any farther for lack of space, but I do mean to make explicit my connection with Sellar's perspective on the development of science (and more importantly scientific modelling practices).

[46] This may seem confusing, but to grasp the idea think of a hypothetical model *M* of combustible material *Q*. Here the target system of the model will have the properties of being a material *Q*, of burning, of releasing heat, of releasing light etc. Now, if these properties are located in a logical space that permits of a huge range of possible states or characteristics of such properties, then it will be far more difficult for *M* to describe the logically possible configurations of the properties or relations of the target system *Q*. It may be the case, for instance, that it is possible for the combustible material to be hot when it does not release light on one occasion, and hot and bright on another occasion. And any explanation based on *M*, therefore, will have a more difficult job describing the logically possible configurations of the properties or relations of the target system, because there will be a great many logically possible configurations to describe given the prescriptions of *M* about the nature of *Q*.

to describe will be very high. To be clear, I am suggesting that when a model *begins* from the location of items in a logical space that permits of a high number of logical possibilities, this negatively affects the standing of the explanatory utility of that model. In this way, a model that *begins* from the location of items in a logical space that permits of fewer logical possibilities will have a greater explanatory utility. Consequently, when we compare two models, then the one that *begins* from the location of items in a logical space that permits of fewer logical possibilities will have a greater explanatory utility.

Thus, my argument is that the measurement that a model *begins* from is directly correlated to the explanatory utility that very same model has. So, if we say that all models $m_1 \dots m_n$ *begin* from logical spaces $l_1 \dots l_n$, where the number of possible states of affairs logical spaces permit is given as $l_1{}^j \dots l_n{}^j$, then the model and logical space pairing $a = \langle m, l_1{}^j \rangle$ has a greater explanatory utility than the model and logical space pairing $a = \langle m, l_2{}^j \rangle$, iff $l_1{}^j < l_2{}^j$, where $<$ says that the number of possible states of affairs logical space $l_1{}^j$ permits is less than the number of possible states of affairs logical space $l_2{}^j$ permits. Likewise, if we have two models m_1 and m_2, then m_1 will have a greater explanatory utility than m_2 iff $a = \langle m_1, l_1{}^j \rangle < a = \langle m_2, l_2{}^j \rangle$, where, again, $<$ says that the number of possible states of affairs logical space $l_1{}^j$ permits is less than the number of possible states of affairs logical space $l_2{}^j$ permits.

But this of course leaves open the most important two questions; namely: What is it that determines how many possible states of affairs a logical space permits of? And: How can we evaluate which measurements (i.e. location of properties of the target system in logical space ([MEAS])) modelling should *begin* from?

I have said that a model that *begins* from a measurement that locates the properties of a target system in a logical space with a lower number of logical possibilities will have a higher explanatory utility. But it is important to recognise that this only applies where it can be shown that modelling is justified in *beginning* from a measurement that locates the properties of a target system in a logical space in the first place. This justification is essential, or else it can be argued that a model that *begins* from a measurement that locates the properties of a target system in a logical space that permits of only one logical possibility – i.e. everything is caused by God – is the model with the best explanatory utility possible.

Now, it seems that any justification for *beginning* from one measurement instead of another cannot issue from intuition, feeling, or fiat. This must be the case or else we find that models that *begin* from measurements that locates the properties of a target system in logical spaces in which every possible state of affairs is caused by God will always have the best explanatory utility. So whilst it is true that modelling can, in principle, *begin* from a measurement that locates the properties of a target system in any random logical space, models that do *begin* from such measurements will not have a high explanatory utility unless their *beginning* from such measurements can be justified. That is, unless a model's *beginning* from a measurement that locates the properties of a target system in logical space l_1^j is justified given the fact that the model could have *begun* from a measurement that locates the properties of a target system in logical spaces l_2^k ... l_n^n. But how, then, could we justify *beginning* model construction with a measurement that locates the properties of a target system in logical space $l_1^{\,j}$ and not a measurement that locates the properties of a target system in logical spaces $l_2^{\,k}$ ... $l_n^{\,n}$?

This is where the importance of [SCI] above becomes apparent. Because [SCI] accounts for the way that science – as a collection of institutional norms and error filters – really does regulate how many possible states of affairs a logical space should permits of. Thus, [SCI] also stipulates which logical spaces a *scientific* model should *begin* from. [SCI] does this by stating that any model that *begins* from a logical space that is powerless to reject possible states of affairs that contravene the dictates of scientific institutional norms and error filters should be taken to have little or no explanatory utility[47] So, according to ([SCI]), any model that *began* from a measurement that locates the properties of a target system in a logical space that contravenes the dictates of scientific institutional norms and error filters would not count as scientific, because that model would not be subject to the institutional norms and error filters that ensure models are, for instance, empirically testable and thus verifiable. And so models are only able to function as the representation-device for *scientific* representations if the measurement from which such models *begin* locates the properties of a target system in a logical space that permits only of possible states of affairs consistent with the institutional norms and error filters of science.

In this way, we see that models are used to explain, and that the explanatory utility of a model is indexed to; (1) a measurement that makes

[47] I ask those who want to say that scientific institutional norms and error filters are also partly based upon fiat to recall what I said in footnote 20 above.

possible (*begins*) model construction in which the properties of a target system are located in logical space; and (2) the justification of locating the properties of a target system in any given logical space in the first place. In the case of (2) here, science ([SCI]) – as the collection of institutional norms and error filters – stipulates which logical spaces a *scientific* model can and cannot *begin* from; i.e. in which logical spaces properties of a target system can be located ([MEAS]) to make possible model construction.

Models, Logical Space, and Content

It is helpful to once again summarise what has been said thus far in this chapter. Firstly, I have argued that models model by means of selective resemblance. I have proposed that the constructive processes that enable selective resemblance are abstraction and idealisation. And, moreover, that the guiding purpose of model construction is explanatory utility. Then I argued that all models are processed artefacts that *begin* from measurement, where measurement is the location of the properties of a given target system in logical space. And this led me to assert that the explanatory utility of a model is directly correlated to: (1) the measurement that the construction of that model *begins* from; and (2) the logical space that measurement operates on.

Given what I have said above, we may want to say that the *content* of the model is the properties of the target system that the model selectively resembles. This content, we may think, is the end product of the constructive processes of abstraction and idealisation, which have each in their own way simplified and distorted the properties of the actual target system. And, what's more, we may think that a model is only able to have this content as a consequence of a measurement ([MEAS]) that located the properties of the target system in some logical space in the act of modelling. In the case of the Erasmus Bridge model in NEMO, then, we may say that the content of the model are some of the properties of the actual Erasmus Bridge in Rotterdam that have been selected to be resembled in the modelling process. And hence that the model as a processed artefact is able to selectively resemble these properties as a consequence of the location of these properties in the logical space of, say, bascule bridge mechanisms.

But surely this is just to say that the content of a model is the properties that *have been located* in a logical space of some kind by the act of measurement. So it follows that the content of a model are the properties *located* in a logical space. In this way, it seems to be the case that the logical space contains the content of the model. Model construction, therefore, can be understood as a kind of reification of a logical space that has been imbued with the properties of a target system.[48] The logical space itself contains the possible states, characteristics, and configurations of the properties or relations of the target system, and the model the processed artefact that is a making objective, in some sense, of these possible states, characteristics, and configurations of the properties or relations of the

[48] Reification is not quite the right term here, because the properties located in logical space need not be given concrete form in the process of model construction. This is evident if we think of mathematic or theoretical models, for instance. To get around this problem of expression, I will coin my own phrase in chapter 4 – objectification. For now, however, I ask the reader to grasp what is intuitively meant be 'reification' here.

target system in the form of a physical, abstract, or fictional thing (a representation-device). A model, in short, is an artefact that makes objective or presents the mapped contents of a logical space. And so the content of a model is equivalent to the contents of a logical space that have been fixed by means of measurement ([MEAS]).

But what, then, causes logical space? Van Fraassen says that the logical space is provided by the background theory in which model construction is couched, and that it permits use to model 'a range of possible states or characteristics' of properties of target systems (Van Fraassen, 2008: 164). I do not want to restrict myself to saying that it is only a theory that provides the logical space by which measurement proceeds. I do not want to say this because a theory, on my [REP] conceptualisation of representation, is simply another representation-device in exactly the same sense as a model. And so I think that a theory can also be understood as an artefact that makes objective or presents the mapped content of a logical space. But then we are left to ask: What does give rise to the logical space in which the properties of a target system can be mapped and located in the act of measurement ([MEAS])? And what, to repeat, is a logical space? Let's begin with the second of these questions.[49]

[49] It may be helpful here to give a brief list of some examples of logical spaces in various branches of science. For example, in physics we have the logical spaces of Newtonian mechanics and relativistic mechanics; in biology we have the logical spaces of evolution or genetics; in sociology we have the logical spaces of law, literature, culture, and gender; and in geoscience we have the logical spaces of geography, geology, ecology, and glaciology. As will be clear, this list is not in any sense exhaustive. There are many more logical spaces relevant to the sciences and, what's more, the boundaries between different logical spaces are often fluid and/or difficult to demarcate (consider the differences between meteorology and climatology, for instance).

One option is to say that logical space is the realm of all possible states of affairs, *à la* Wittgenstein (c.f. Wittgenstein, 1922: TLC 1.13). Although there has been a some disagreement amongst philosophers as to how we should conceptualise logical space on Wittgenstein's account, for heuristic purposes we can think of Wittgenstein's conception of logical space as a collection of 'possible worlds, each distinct set of which constitutes a distinct logical place' (Stenius, 1960: 55).

Another option is to think about logical space as a logical space of reasons, in Sellars (1956: § 36) sense. Sellars argued that:

> The essential point is that in characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.

On Sellars view, then, logical space is a region in which one can operate once one is 'at home with normative discourse, responsive to reasons as such, and sensitive to standards of correctness and appropriateness which are applied to one's own states and those of one's comrades' (deVries, 2011).

Yet another option for conceptualising logical space is to follow Rayo (2013) and argue that it is the set of distinctions amongst the way the world is to be. In this way, logical space a set of 'live options' when we theorise about how the world could be (Rayo, 2013: 3) What's more, on Rayo's view, one can 'shape' logical space by accepting or rejecting 'just-is' statement, such as, "Water just is H_2O" (Ibid.: e.g. 18).

Now, it is clear then that the question as to what logical space *is* remains a live debate in philosophy, and I do not expect to completely resolve that debate here. Instead, I want to bring this discussion back to my discussion about scientific models by returning to the question: What gives rise to the logical space in which the properties of a target system can be mapped and located in the act of measurement ([MEAS])?

Having quickly reviewed Wittgenstein's, Sellar's, and Rayo's conception of logical space it becomes clear that it will be all but impossible to give a conceptualisation of logical space that does not share one common feature: logical space is always conceptualised as something reasoning – i.e. intentional or semantic – agents use to think about the world. In this sense, it may be straightforward to argue that logical space is merely 'about' the world and not 'in' the world. But if we think back to my rejection of the 'in'/'about' dichotomy in the case of representations above ([REP]), and the anti-dualist stance of this essay in general, then it is unsurprising that I want to advance my own conceptualisation of logical space. That conceptualisation, stated in the broadest possible terms, is as follows:

[LOGSP] Logical space is the space in which the world and agent's activities coincide.[50]

Where by 'coincide' I mean something like co-occur, overlap, interact, or co-exist.

---

[50] As far as I know there is no conceptualisation of logical space that would contradict [LOGSP], but even if there were it is difficult to imagine how it could do so without implicitly endorsing dualism. I will take this chance, therefore, to once again state that I reject dualism in all its varieties. I take this to be sufficient to defend [LOGSP].

Given this conceptualisation of logical space, my answer to the question, what instantiates the logical space in which the properties of a target system can be located and subsequently modelled?, becomes clear: I hold that:

[LOGSP*] Logical space is jointly constituted and jointly instantiated by agent(s) and world. Logical space emerges from the relation between agent and world.[51]

In this way, I hold that logical space – like representation ([REP]) – is 'in' and 'about' the world at the same time.[52]

Now, let us return to our discussion about the content of a model. I said above that a model is an artefact that makes objective or presents the mapped contents of a logical space. And armed with my [LOGSP*] conceptualisation what instantiates logical space, it follows that if something is located or mapped in a logical space it has been jointly constituted and instantiated by agent(s) and the world. The content of a model is then, on my view, constituted and instantiated by agent *and* world in a logical space which can be presented in the process of modelling; i.e.

[51] It has been suggested to me that [LOGSP*] seems to imply that logical space is a dyadic predicate with agent and world variables. This description of logical space at the level of language would be correct were we use the term "logical space" to refer to that thing that is jointly constituted and jointly instantiated by agent(s) and world. However, logical space also has a relationship with the processed artefacts that are representation-devices (i.e. scientific models). If anything, then, [LOGSP*] must be understand as a triadic predicate in *exactly* the same sense as representation [REP].

[52] It is important to note that there can be different logical spaces in the sense that the logical space of Newtonian mechanics differs from the logical space of Relativistic mechanics. Still, I take these differences to obtain at a higher level of generality than what is captured by my [LOGSP*] definition, and so I hold that all logical spaces can be subsumed under the conceptualisation [LOGSP*].

the process of creating an artefact that makes objective or presents the mapped contents of a logical space.

The implication of this view is that I must qualify what was meant throughout this chapter when I have said that models selectively resemble their target systems, and that abstraction and idealisation are the constructive processes by which models are made to resemble their target systems. In both of these cases it is now apparent that what is meant by target systems is, in fact, target systems *as located or mapped in logical space*. I have just said that logical space is jointly constituted and instantiated by agent(s) and the world so it is not the case that what models selectively resemble is completely cut off from the world. But it is case that what models selectively resemble – i.e. the content of models – is also bound to the activities of agents.

I assume that one criticism of the position I am now elaborating will focus on my claim that modelling must *begin* from a measurement, where a measurement is a location of the properties of a target system in logical space ([MEAS]). I assume that the criticism will be that the content of models, even if it is located or mapped in logical space, is first and foremost located or mapped in logical space by agents who encounter target systems in their experiences of the world. So ultimately the content of models, as this criticism goes, is more the result of agential factors than worldly factors. My response is simple: in order for the agent to locate properties of target systems in logical space there must be a target system with properties to locate. And, moreover, according to my definition of logical space it is only in the coincidence of agent (locator) and target system (located) that we have the constitution and instantiation of the logical space in which the properties of the target system are located in the first place. So any dualistic claim that attempts to prioritise an agential or worldly account of the content of models will fail as a result of my [LOGSP*] conceptualisation of logical space.

What's Special about Scientific Representations, again?

I said in the first sub-section of this chapter that by accepting [SCI] I am committing myself to the claim that scientific representations are special, because they aim at giving the most approximately accurate account of reality available. By evaluating all that I have said in this chapter, I want now to defend this claim explicitly.

Recall first of all my conceptualisation of representation in chapter 1:

[REP] Representation is the structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented.

Scientific representations will be a manifestation of the same structural relationship, with the added feature of being aimed at giving the most approximately accurate account of reality available. But how, then, can this feature actually be added to [REP]? How, in other words, can we account for the claim that scientific representations *are* a structural relationship that aims to give the most approximately accurate account of reality available?

To start we can make explicit once again the fact that the representation-device of all scientific representations will be a specific kind of processed artefact: a scientific model. Then we can restate the difference between *scientific* models and other models as I spelled it out in the earlier sub-section on models and measurement. I said, referring back to my [SCI] conceptualisation of science above (pg. 86), that:

> [SCI] […] stipulates which logical spaces a *scientific* model should *begin* from. [SCI] does this by stating that any model that *begins* from a logical space that is powerless to reject possible states of affairs that contravene the dictates of scientific institutional norms and error filters should be taken to have little or no explanatory utility.

Now, we have seen in the previous sub-section that when a model *begins* from a measurement that has located the properties of a target system in logical space, the model is in fact presenting the possible states, characteristics, and configurations of the properties of the target system that are *contained* in the logical space. So *scientific* models are scientific in virtue of only presenting the possible states, characteristics, and configurations of the properties of the target system that are *contained* in logical spaces *that do not contravene the dictates of scientific institutional norms and error filter*. Put briefly, *scientific* models are scientific in virtue of presenting only the possible states, characteristics, and configurations of scientifically admissible logical spaces. So scientific models are a special kind of processed artefact that functions as a representation-device.

Above, I conceptualised a logical space as jointly constituted and instantiated by agent and target system (world). And we have just seen that *scientific* models are scientific in virtue of presenting only scientifically admissible logical spaces. Here, then, we have a relationship between *scientific* model and the *scientifically admissible* logical spaces that are jointly constituted and instantiated by agent and target system (world). So we have a special kind of relationship between the agent and world (scientifically admissible logical space), and the representation-device (a scientific model that is a processed artefact that makes objective or presents the mapped contents of scientifically admissible logical spaces). What is special about this relationship is that it is formed by scientific admissibility in the sense of [SCI]. This special relationship, I propose, *is* scientific representation. In other words, when models are scientific in virtue of presenting only scientifically admissible logical spaces, and all logical spaces are taken to be constituted and instantiated by agent and world equally, then we are able to have a particular kind of structural relationship between *scientific* model and *scientifically admissible* logical space. So:

[SCIREP] A scientific representation is the objective structural relationship between scientifically admissible logical spaces in the sense of [LOGSP*] and a representation-device in the sense of a scientific model.[53]

I would also endorse a reformulation of [REP] in these terms, so that it becomes:

[REP*] Representation is the objective structural relationship between logical spaces in the sense of ([LOGSP*]) and a representation-device of any kind.

[53] To repeat what I said at the end of chapter 1, the relationship is structural because all components are required if we are to have a case of scientific representation – all components mutually support one another and there would be no scientific representation without a relationship between all three (two if you count logical space as one component jointly constituted and instantiated by both agent and world).

And with this formulation of [SCIREP], and this reformulation of [REP], we can finally see why scientific representations are of a qualitatively different kind than other representations. The reason is because the logical spaces that make possible the structural relationship that *is* scientific representation will be 'scientifically admissible', and subsequently make possible *scientific* representation-devices – i.e. *scientific* models as processed artefacts that present only the scientifically admissible logical spaces. Subsequently, when scientifically admissible logical spaces enter into a structural relationship with the scientific models they beget, a scientific representation is enacted ([SCIREP]).

To be scientifically admissible a logical space must permit only of possible states of affairs consistent with the institutional norms and error filters of science as set out by [SCI]. So a logical space that permitted of, for example, a geocentric solar system would not be scientifically admissible. Nor would any logical space that permitted of states of affairs involving phlogiston, luminiferous aether, phrenological explanation, a static universe, earth being flat or being five thousand years old etc. Scientific admissibility, therefore, is the mechanism that makes explicit which logical spaces are to be preferred and which logical spaces are to be rejected altogether in the process of modelling.

And how, then, can we justify the notion of scientific admissibility that is packed into the structural relationship that *is* scientific representations? It is justified, I argue, because scientific admissibility is fixed by institutional norms and error filters that, according to [SCI], are aimed at discovering the objective character of the world. To reiterate what I have been saying then: it is the mechanism of scientific admissibility that allows for scientifically admissible logical spaces, scientific models, and thus the structural relationship between scientifically admissible logical spaces and scientific models that *is* scientific representations ([SCIREP]). Moreover, it is the mechanism of scientific admissibility that ultimately differentiates scientific representations from other kinds of representation. Why? Because scientific representations are special because they are formed by the mechanism of scientific admissibility, and – according to [SCI] – it is the mechanism of scientific admissibility that makes it the case that scientific representations, in contrast to other kinds of representation, *aim* to give the nearest approximation of reality possible.

Chapter 3.

Information

The purpose of this chapter is to motivate and formulate a novel conceptualisation of information. To do this, I first consider various competing conceptualisation of information, and draw out the fundamental – in fact, pre-theoretical – disagreement as to whether information should be conceptualised as *substantive/extensionalist* ('in' reality) or *semantic/intentionalist* ('about' reality). I then consider philosophical attempts to move beyond this dichotomy between the *substantive/extensionalist* and the *semantic/intentionalist* conceptualisation of information, and find that the dichotomy has yet to be overcome. Finally, I consider what any conceptualisation of information should account for, before formulating my own novel conceptualisation of what information *is*.

Conceptions of Information

Information is a notoriously problematic concept. Information was described by Floridi (2011: 81) as a 'polymorphic phenomenon and a poly-semantic concept', which 'as an *explicandum*, can be associated with several explanations depending on […] the cluster requirements and desiderata of a theory'. In a similar vein, Shannon (1993) argued that:

> The word ‘information’ has been given different meanings by various writers in the general field of information theory. It is likely that at least a number of these will prove sufficiently useful in certain applications to deserve further study and permanent recognition. *It is hardly to be expected that a single concept of information would satisfactorily account for the numerous possible applications of this general field.* (Shannon, 1993: 180; italics added)

Information, then, as both a phenomenon and a concept, is equivocal. And it is this equivocity that has been perhaps the central and recurring theme in the surveys of information undertaken over the last thirty years.[54] To make a start on analysing the various conceptions of information, therefore, is no easy task. The best option, I think, is to find some point of dispute to set about grouping – and hence differentiating – various conceptions of information. I will turn to that task after canvassing one distinction that will be integral to what follows in this chapter.

In anticipation of what is to come it is important to make clear the simple distinction between information and the information-carrier/bearer that has been prevalent in discussions of information. What this distinction amounts to is the claim that information differs from that which carries, contains, or bears information. A simple example can serve to make this clear. A newspaper, full with news stories, advertisements, and obituaries, seems to carry, contain, or bear information. But, according to certain theorists, we would not want to say that the

---

[54] C.f. Debons and Cameron, 1975; Larson and Debons, 1983; Machlup and Mansfield, 1983; Braman, 1989; and Losee, 1997. In what follows I will continually refer to perhaps the most recent survey of information: *The Handbook of Information* edited by Adriaans and van Benthem, 2008.

newspaper – taken as a collection of pieces of paper marked with ink – *is* information. Rather, we may feel inclined to think that the information is, for example, the news of economic turmoil in a particular nation state, the announcement of discount price electronic goods, or the tragic news of someone's death. Thus, it has been held that information – whatever it is – should be differentiated from the information-bearer. That is not to say, however, that this distinction is always straightforwardly identifiable, as will become clear in the conceptions of information I will now consider. For my part, I will refrain from commenting on the information/information-bearer distinction until it becomes necessary at the beginning of the next chapter.

In their (2008), Adriaans and van Benthem argue that there are three major concepts of information.[55] They present these concepts as follows:

Information-A	Knowledge, logic, what is conveyed in informative answers
Information-B	Probabilistic, information-theoretic, measured quantitatively

[55] In reaching this point they already delicately avoid views on information that take it to be 'just a metaphor' and views on information that suppose that we should 'chart' the position of information in regards to related concepts (Adriaans and Benthem, 2008: 10). I too will delicately avoid the first of these views by not mentioning it at all. But this entire essay should be seen as an attempt to put the second view into practice by specifying the relation between information and one other concept in particular: (scientific) representation.

Information-C	Algorithmic, code compression, measured quantitatively[56]

(Adriaans and Benthem, 2008: 11)

By taking Adriaans and van Benthem's demarcation as a starting point, I will now consider the features of more specific conceptualisations of information.[57] Let's begin with Information-B, on their demarcation. The paradigmatic example of a probabilistic account of information is Shannon's (1948) work on 'channel transmission', which is considered by many to be *the* basis for information-theoretic approaches to information. Shannon developed a mathematical theory of the codification and communication of data and signals, which is now widely referred to as the communicative theory approach to information (c.f. Shannon and Weaver, 1949; Adriaans and van Benthem, 2008, especially Harremoes and Topsøe, 2008). The central idea of Shannon's approach was to define information in terms of probability space distributions, where 'the amount of information in a message is the negative base-2 logarithm of the probability of its occurrence from a given source over a given channel […] measured in 'bits', which has become a household term' (Adriaans and van Benthem, 2008). In this way, the information that is 'received' by a 'receiver' is measured in terms of expected reduction of uncertainty (i.e. probability) (Ibid.).

[56] It is worth noting that Adriaans and van Benthem accept that they are 'oversimplifying a bit', but this does not detract from the heuristic value of their demarcation (2008: 11).
[57] I will only give a very cursory introduction to the different specific conceptualisations of information, but this will be sufficient to substantiate my later argument about the underlying tensions between the current conceptualisations of information.

Shannon's idea, therefore, was that information is always already *about* something, and as such a description of information should be distinguished from the something that it is about. In the terms introduced above, Shannon's idea was that a description of information can be given by analysing the content of an information-bearer, where the content is said to be composed of bits of data. This allowed for Shannon to develop a technical, quantitative description of information, which, in his words, allows for a 'reproducing at one point either exactly or approximately a message selected at another point' (Shannon, 1948: 1). By borrowing from Gibbs (1906) development of Boltzmann's formula, $S = k \ \log W$, in physics, Shannon equated the amount of distributions in probability space given by a system of messages A with the 'communication of entropy' of a system of messages A. This was formalised as:

$$H(Pr) = -\Sigma_{(x \in A)} Pr(x) \log_2 Pr(x)$$

Where $Pr(x)$ give the probability of message x in A. In this way, Shannon tells us, 'the amount of information I in an individual message x is given by':

$$I(x) = -\log Pr(x)$$

So a message x has a specific probability p_x between 1 and 0 of occurring, so if $p_x = 1$, then $I(x) = 0$: in other words, 'if we are certain to get a message it contains literally no news at all' (Adriaans, 2012). Interestingly, this also means that if we have two unrelated messages, say, x and y, then $\mathrm{I}(\mathrm{x}\ \text{and}\ \mathrm{y}) = \mathrm{I}(\mathrm{x}) + \mathrm{I}(\mathrm{y})$. So information, on Shannon's characterisation, like entropy on Boltzmann's characterisation, is extensive: it is an additive property so that the value of the property described (information, entropy) is proportional to the size of the system described.

Now, Shannon's theory is considered to have been a great success and still forms the foundational basis of much of contemporary information-theory today. Still, Shannon's theory unapologetically defines information exclusively in terms of the code length of a message measured in units of data – i.e. bits. Consequently, Shannon's approach to information is not able to account for the propositional content of informational messages. Consider, for example, the two messages, "Utrecht is in China", and, "The earth is made of chocolate". In the case of these two message, Shannon's conceptualisation of information holds that both carry the same amount of information even if their meaning – or propositional content – is seemingly quite different. This, it seems, is the downside of adopting a probabilistic approach to information, because the probability space that is central to Shannon's conceptualisation only makes room for a consideration of information via the encoded data content of messages, and this alone will not be enough to tie the information 'carried' by a message to the meanings of such messages themselves (Shannon, 1948: 1). The probabilistic conceptualisation of information, in other words, concerns only the substantial, data content of information (bits) and not the semantic, propositional content of an information-bearer. That is, it defines the information carried by the information-bearer in purely bit-data terms.

Returning to Adriaans and Benthem's demarcation of the three major concepts of information, I will now consider Information-C conceptualisations of information. The paradigmatic example of an

Information-C conceptualisations of information is the algorithmic or computational account of information known as Kolmogorov complexity[58]. Kolmogorov's approach to the conceptualisation of information overlapped with Solomonoff's attempts to universalise Carnap's method of 'assign[ing] a priori probabilities to any possible string of symbols that might represent the universe' (Solomonoff, 1997: 6). Solomonoff's *universal distribution* stated that we should:

> consider the set of all possible finite strings to be programs for a universal Turing machine U and define the probability of a string x of symbols in terms of the length of the shortest program p that outputs x on U. (Solomonoff, 1960, 1964a, 1964b).

For Kolmogorov, the length of a program p required to output x on U is the shortest code of the program p. And, in this way, we can correlate the shortest – or optimal – code of a given program p to the complexity of the string x. So the complexity of a string x is defined as the length of the shortest code of a program p that computes x on some fixed universal Turing machine. Thus, so long as we work with an enumerable set of prefix-free programs, we can easily give the probability distribution for all Turing machine U computable strings based upon their relative complexity.[59] This, according to Kolgomorov, allows us to arrive at a measurement of information in terms of the shortest code – or smallest computer program – that can compute or calculate a given string on a universal Turing machine. In this way, Kolmogorov complexity 'starts with the shortest code as fundamental and derives an *a priori* probability

---

[58] See Chaitin (1987) for a more recent presentation of Kolmogorov's position.
[59] For more information about universal Turing machines see Herken, 1992.

from it' (Adriaans and van Benthem, 2008: 12). And information *is* the complexity of a string as determined by the length of the code required to compute the sting in question.

To be clear, then, the algorithmic or computational account of information measures *and* defines information in terms of the complexity of any possible finite string of symbols, where the probability of a string can be derived from its complexity. It follows, therefore, that 'highly regular strings will have low complexity, while highly random strings have high complexity' (Ibid.: 6). This particular feature of the algorithmic or computational account of information has turned out to be of great use to philosophers, information-theorists, and computer scientists alike, because it can be argued to provide us with a general theory of induction (c.f. Grünwald 2007); to allow us to formulate probabilities and informational content for individual objects, even including the natural numbers (c.f. Chaitin, 1987); and to 'allow us to formulate an objective a priori measure of the predictive value of a theory in terms of its randomness deficiency: i.e., the best theory is the shortest theory that makes the data look random conditional to the theory' (c.f. Vereshchagin and Vitányi, 2004) (Adriaans, 2012).

It is not difficult to see, then, why the algorithmic or computational account of information has become one of the pre-eminent contemporary theories of information. Philosophers, however, can still be critical of any algorithmic or computational conceptualisation of information. The reason for this is clear: because the algorithmic or computational conceptualisation of information – similarly to Shannon's probabilistic conceptualisation – defines information exclusively in terms of the computational complexity of a string (or message) and says nothing about the propositional content of that string (or message). This is the case because the information of a string (or, again, message) is taken to be equivalent to the size in bits of the smallest computer program – or code – required to calculate or compute the string (or, once again, message). Thus, information is reduced to informational content conceived as a mere data (i.e. bits). And this provides no way to differentiate the meanings attached to the various strings (or messages) (consider once again the messages about Utrecht and chocolate above).[60] Once again, then, we find that the substantial, data content of information (bits) and not the semantic, propositional content takes centre stage. Information is assumed to be conceptualisable solely in terms of the complexity of the bit-data carried by an information-bearer (strings/messages), leaving aside all issues regarding the meaning of bit-data altogether.

Not surprisingly, the last of Adriaans and van Benthem's concepts of information that I will consider is somewhat different to the two concepts I have now discussed. Information-A conceptualisations of information are concerned with the 'communication-oriented sense of information' (Adriaans and vam Benthem, 2008: 6). That is, with the study of those informational phenomena that are related to semantic meaning, knowledge, and other such concepts found in the domain of human activity. For the most part, 'communication-oriented' approaches to information are studied within the confines of logic and linguistics. Logical accounts of, for example, inferential processes and belief states are

[60] Consider two 'strings', "Utrecht is in China", and, "The earth is jelly!", which will be computable given the same length code of a computer program, but do appear to quite different propositionally speaking.

informational in the sense that they consider the ways that 'agent acquire new information about what the real world is like, through acts of observation, linguistic communication, or deduction' (Ibid.: 11). Subsequently, Adriaans and van Benthem stress that there are three vital features of the 'communication-oriented' approach to information: (1) agents who, in their terms, 'represent' and use information; (2) the dynamic nature of information change; and (3) the fact that information is always 'about' some aspect of the world – i.e. information is always meaningful (Ibid.).

Now, the 'communication-oriented' approach can be further decomposed into a number of different sub-approaches. But I think it is possible to group these various sub-approaches together in accordance with a commitment endorsed by all communication-oriented approaches to information; namely, the commitment to the claim that information is *semantic* or that information has *semantic content*. In view of this commitment, I agree with Floridi (2011) that *all* communication-oriented approaches to information seek to provide answers to questions such as: "How can something count as information? And why?", "how is information related to error, truth and knowledge?", and, "when is information useful?" (Floridi, 2011). What's more, it is through the process of formulating viable answers to these questions that communication-oriented approaches to information connect the concept of information to other relevant forms of doxastic, mental, and epistemic phenomena. For here on, therefore, I will say that all Information-A concepts of information adopt a semantic conceptualisation of information.[61]

The bedrock of all semantic approaches to information is the *general definition of information*. The general definition of information holds that:

[GDI] σ is an instance of information, understood as semantic content, if and only if:

(GDI.1) σ consists of one or more *data*;

(GDI.2) the data in σ are *well-formed*;

(GDI.3) the well-formed data in σ are *meaningful*.

(Floridi, 2011: 84)

[GDI] has been defended by a number of philosophers, computers scientists, and information theorists. For example, Davis and Olson (1985: 200) argue that 'Information is data that has been processed into a form that is meaningful to the recipient'. Checkland and Scholes (1993: 303) argue that 'Information equals data plus meaning'. Warner (1996: 1) argues that 'data will need to be interpreted or manipulated [to] become information'. And Quine (1970: 3-6, 98-99) proposes that there is an equivalence relation between 'likeness of meaning', 'sameness of proposition', and 'sameness of objective information'.

[61] We should recognise, however, that a probabilistic – Information-B – semantic approach to information was developed by Bar-Hillel and Carnap (1953). Bar-Hillel and Carnap began from Shannon's communicative theory of information, but sought to replace Shannon's quantitative notion with a qualitative logic of information spaces. Bar-Hillel and Carnap's work will occupy my thoughts later in this chapter, when I come to draw out the underlying tensions between competing conceptualisations of information.

(GDI.2) stipulates that in order to be meaningful, data must first be well-formed. In other words, data must first be 'clustered together correctly, according to rules (syntax) that govern the chosen system, code or language being analysed' (Floridi, 2011: 84). The term syntax here must be understood to encapsulate not just linguistic clustering, but rather as the general forming, composing, or structuring of any σ which is to count as an instance of information. As Floridi (Ibid.) remarks, 'Engineers, film directors, painters, chess, and gardeners speak of syntax in this broad sense'.

(GDI.3) stipulates that in order to be information, well-ordered data must also be meaningful. In other words, the data in question must comply with the semantics of the chosen system, code or language in question. Once again this should not be understood solely in terms of linguistic meaning. A map, for instance, can be 'visually' meaningful for a reader, and so too can an exhibit in a science museum, like NEMO. What is important, therefore, is only that information is used by agents and is, in some sense, about the world.

Now, it will be evident given [GDI]'s emphasis on syntax and semantics that logic has played a central role in the investigation and articulation of semantic approaches to information. Thus, by taking my cue from van Benthem and Martinez's (2008) survey of logical perspectives on information, I will now briefly outline some examples of the semantic conceptualisations of information. Each of these conceptualisations, I think, is concerned in their own way with the question: "What is conveyed in informative answers?", and can be subsumed under the Information-A conception of information currently under discussion.

One semantic conceptualisation of information that can be logically articulated is the conception of information as a range (van Benthem and Martinez, 2008: 217). The basic idea here is that 'The greater one's range of options for what the real world is like, the less information one has' (Ibid.). Thus, the information as range conceptualisation of semantic information holds that the more uncertain you are about how the world is, the less information you have. It follows, then, the logical articulation of information as range will be modal, because information is defined in terms of modal space and in/consistency. This articulation could potentially be given in either a doxastic, epistemic, alethic, deontic, or temporal modal logic, and the result would be the same. What is important, then, is only that any instance of information σ is definable as the set of possible worlds or, perhaps, the set of possible states of the universe, excluded by σ. Thus, an instance of information is conceptualised as encoded in the range of models that determines what should be discounted from consideration. Information, therefore, is the range of possible options for how the situation actually is, such that the more information we have, the fewer possible options there are available.

The answer that the information as range conceptualisation of semantic information gives to the question, "What is conveyed in informative answers?", is that what is conveyed in informative answers is a particular set of possible ways something could be. In other words, what makes an answer informative is what possible worlds it admits of and what possible worlds it excludes. And so, crucially, in order for an informative answer to be informative, agents must be able to make sense of what the answer says in regards to admitted and excluded possible worlds. This is just to say that the answer must be meaningful within a certain system, code, or language that pertains to possible worlds. Given, then, that a semantics for modal logic can be defined using possible worlds, it becomes clear why a modal logical articulation of information is compatible with the information as range conceptualisation of semantic information. The reason, to state it clearly, is because any modal logic can be interpreted as a language that codifies semantic information in terms of the model spaces relevant to, for example, belief (doxastic logic), knowledge (epistemic logic), or time (temporal logic) (c.f. Kripke e.g. 1959, 1963). Thus, according to the information as range conceptualisation of semantic information, all model logical calculi define what information *is* by codifying possible states of affairs.

Another semantic conceptualisation of information that can be logically articulated is the idea of information as correlation (van Benthem and Martinez, 2008: 217). Floridi (2011) calls the idea of information as correlation the 'systemic approach to information'. The central notion of the systemic approach to information is that information should be conceptualised as being about something relevant to agents, and hence should involve connections between different systems. In contrast to the information as range approach, then, the systemic approach supposes that that any instance of information σ is definable *not* through an *a priori* calculus of possible states (modal logic), but rather through the factual content that σ carries with respect to a given situation. So information tracks possible transitions in a system's – or situation's – states space under normal conditions by means of a logic of some form.[62] And so, as Martinez (20014) put it, the systemic approach to information concerns 'the *situatedness of information*; that is, its dependency on the particular setting on which an informational signal occurs'. A good examples of the situatedness of informaiton is the starry sky in which the same pattern of stars, at different moments in time and locations in space, will in general convey different information about one's own or another person's location in space and time.

The answer that the information as correlation (systemic) conceptualisation of semantic information gives to the question, "What is conveyed in informative answers?", is that what is conveyed in informative answers is factual content about a given situation. Thus, we see again that for an informative answer to be informative at all requires that it is taken to be about some situation by an agent. That is, it requires that the answer is meaningful within a certain system, code, or language. For the most part, systemic approaches to information adopt a form of 'situation semantics' in which 'meaning arises from the interaction of organisms and their information-rich environment' (Ibid.: 238). This 'situation semantics' is buttressed by 'situation theory' more generally (Barwise and Perry, 1983), which is standardly holds that:

> everything is part of a structured reality which is full of uniformities. Organisms are 'attuned' to those regularities, and that allows them

[62] For "state space" read a space that is able to represent(!) the different states of a system. This state space could also be dynamic; i.e. phase spaces. For examples of the logical framework of system approaches see Barwise, 1988; Barwise and Seligman, 1997; or Devlin, 2006.

> to survive. Information is a pervasive aspect of reality, prior to cognitive action (Benthem and Martinez, 2008: 239).

It follows, therefore, that systemic approaches to information endorse a form of semantic externalism, which postulates that information is already immanent in the environment in which agents operate. Situation theory was developed to account for exactly this requirement and provides the situational semantics by which the systemic approach to information is able to get off the ground.[63] But this is certainly not to say that situation theory does not face issues of its own. The most pressing of these issues is that there is currently no consensus as to what exactly a "situation" is (c.f Kratzer, 2014).

The final semantic conceptualisation of information I will consider is what I term the veridical approach to semantic information (c.f. Floridi, e.g. 2004c). On this approach, semantic information encapsulates truth, exactly as knowledge does, and so the general definition of information [GDI] takes on an additional clause – call it (GDI.4) – which states that as well as being well-formed and meaningful, data must also be *truthful* to count as information. Opponents of the veridical approach to semantic information will argue for an alethic neutral account of information along the lines of [GDI] above. But the defender of the veridical approach to semantic information will concur with Dretske's claim that:

---

[63] Of particular importance here is the work of Dretske (1981), which I will come to discuss in greater detail in the next sub-section.

> […] *false* information and *mis*-information are not kinds of information—any more than decoy ducks and rubber ducks are kinds of ducks. (Dretske, 1981: 45)

And Grice when he said:

> false information is not an inferior kind of information; it just is not information (Grice, 1989: 371)[64]

To get an idea of the veridical approach to semantic information, we can refer to Floridi's (2011a: 26) 'blueprint of the mechanism that underlies the truthful nature of semantic information'. Floridi says:

> 1. "the beer is in the fridge" qualifies as semantic information if and only if
>
> 2. "the beer is in the fridge" is true; this is the case if and only if
>
> 3. "yes" is the correct answer to (i.e., correctly saturates by correctly verifying and validating) the question "is the beer in the fridge?"; this is the case if and only if
>
> 4. "is the beer in the fridge?" + "yes" generate an adequate model m of the relevant system s; this is the case if and only if
>
> 5. m is a proxy of s and proximal access to m provides distal access to s; and finally this is the case if and only if
>
> 6. reading/writing m enables one to read/write s. (Ibid.)

This, essentially, is to say that some instance of semantic information σ about a system s only qualifies as semantic information in virtue of the fact that an agent in possession of σ is able to access a particular feature

[64] Other defences of a truth-based definition of semantic information can be found in Barwise and Seligman, 1997 and Graham, 1999.

of system *s* as a consequence of a model *m* that is generated by the very fact that a particular feature of system *s* is verifiably there to be accessed.

As complicated as all this may seem, the veridical approach to semantic information simply takes information to be definable in terms of truth, whereby information is not merely 'truth-bearing but truth-constituting', like knowledge, but unlike beliefs or propositions (Ibid.: e.g. 27-28).[65] It follows, then, that the answer that the veridical approach to semantic information gives to the question, "What is conveyed in informative answers?", is that what is conveyed in informative answers is truth or truthfulness. Importing the notion of truth into discussions of information can only make sense on a semantic conceptualisation of truth, however, because the claim that informative answers convey truths requires that the answers in question are meaningful within a certain system, code, or language. Put simply, the veridical account of information must be semantic because truth – whatever that is – seems to be intimately connected to meaning.[66]

I have now presented three semantic approaches to information, which can be broadly categorised under Adriaans and van Benthem's (2008) Information-A concept of information: information as range, the systemic approach, and the veridical approach. Earlier, I also presented the paradigmatic example of their Information-B concept of information

[65] I do not want to go into any further detail *here* about information being truth-constituting. But it will be part of my considerations in the final chapter of this essay to discuss the accuracy of information *and* representation, and in my conclusion to discuss the way that truth connects to information and representation.
[66] The literature on the connection between truth and meaning is substantial to say the least, but I assume that the reader is already quite familiar with the relevant philosophical discussions.

– Shannon's communicative theory of information; and the paradigmatic example of their Information-C concept of information – Kolmogorov complexity. Given this multiplicity of approaches to information, I take Adriaans and van Benthem's demarcation to be a helpful starting point for the task grouping – and hence differentiating – the various conceptions of information available. But I also think that it is possible to draw out a set of more fundamental differences that will serve to keep apart these varying conceptualisations of information. Spelling out these more fundamental differences and considering their validity will be my aim in the next sub-section.

Information: Location and Independence

I noted in my discussion of Information-B and Information-C above that in both cases information was defined probabilistically in terms either of the measurement of the code length of a message (communicative theory of information), or following from a determination of the complexity of a string (Kolmogorov complexity). I also identified that in the case of both Information-B and Information-C the unit of measurement eventually used to determine both the length of a code or the complexity of a string was a 'bit' of data. What was interesting about this was that the Information-B and Information-C conceptualisations of information could not take into account the propositional meanings of different messages or strings, and so were compelled to define information in purely data-based, semantic-less, terms.

In contrast, the semantic conceptualisations of information I considered (i.e. Information-A) all approached information as something meaningful. This took the form of an understanding of how the world could and could not be (possible worlds); an understanding of how a situation is (systemic approaches); or an understanding of what is true (veridical approaches). In each of these cases, then, the definition of information given was bound to the semantic activities of an agent in such a way so as to be about *something* (i.e. possible worlds, situations, or truths).[67] In this way, the definition of information endorsed by all sematic conceptualisations of information is not concerned with the data content of a message or sting, but rather with the way that a message or string is put to use by an agent; i.e. the propositional content of the message or string. The reason for this is simple, because any definition of information in terms of possible worlds, factual content about a situation, or truthfulness requires first that there is an agent with a meaningful way of thinking about possibility, situational facts, or truth.

Now, the two central questions that arise following my introduction of the various conceptualisation of information in the previous sub-section are as follows: Firstly: Where is information located? And, secondly, Can there be information without an informee? Before considering the implicit or explicit answers given to these two questions by the various conceptualisation of information currently under consideration, it is helpful to note that the ontological status of

[67] It should be noted that there is no strict division between these semantic conceptualisations of information. Veridical approaches need not contradict e.g. a possible world approach, because one can just claim that a content specified as excluded by a possible worlds approach is only information if true.

information itself remains obscure. Most people would agree that there is no information without data representation of some kind, whether semantic or not. But the problem with this claim is that it is not clear what is meant by data *representation* in this regard. Some interpret the claim materialistically, whereby it is impossible to have physically disembodied information; i.e. representation is equal to material implementation (Landauer and Bennett, 1985; Landauer, 1987, 1991, 1996). But this interpretation does not seem to square with any conceptualisation of information that permits of abstract data representation, such as sentences or mathematical formalism. Given this state of affairs, another option is to argue that not every representation of data is required to be physically or materially implemented. But this still makes it no clearer what a *representation* of data actually *is*. One further option has been elaborated by Wheeler (1990: 5), who says:

> Otherwise put, every "it" — every particle, every field of force, even the space-time continuum itself—derives its function, its meaning, its very existence (even if in some contexts indirectly) from the apparatus-elicited answers to yes-or-no questions, binary choices, *bits*. "It from bit" symbolizes the idea that every item of the physical world has at bottom—a very deep bottom, in most instances—an immaterial source and explanation; that which we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and that this is a *participatory universe*.

Wheeler's position has gained some support within contemporary physics (c.f. Vedral, 2010, for example). But it is radical in the sense that it proposes that if we conceive of information as data representation, then a representation cannot to be taken to be of the physical or material, but, instead, of an 'immaterial source' of the physical or material. Subsequently, the claim that information is data representation will, on Wheeler's view, be compatible with Wiener's (1954: 132) claim that:

> Information is information, not matter or energy. No materialism which does not admit this can survive at the present day.

Philosophical discussion about the ontological status of information, therefore, is either unclear or has reached the point that information is taken to be an ontological category unto itself. I will say no more about the ontology of information here, but I ask the reader to keep this in mind – along with the distinction between information and information-bearer – as they continue through this chapter.

Returning, then, to the questions: Where is information located?, and, Can there be information without an informee?, we find that the various conceptualisations of information introduced in the previous sub-section are committed to a variety of differing answers.

Take the question, where is information located?, to begin with. Information-B conceptualisations of information will argue that information is to be located 'in' a message occurring from a given source over a given channel, and is measured in bits. The message is simply a discrete unit of communication of one form or another, whether we are talking about a sentence spoken as part of a radio broadcast or one page of a newspaper. And information, therefore, is 'in' this message to the extent that it can be measured in the basic units of data.

We find much the same kind of answer for Information-C conceptualisation of information. In this case, information is the complexity of a string as measured by the shortest code, or smallest computer program, required to compute or calculate the string on a universal Turing machine. Information, then, is located 'in' the string as a property of complexity. This property, of course, can only be measured in the basic unit of information (bits of data, again), but nonetheless it is a property of any possible string that can be apprehended by a process of computation.

In the case of Information-A, however, the response to the question, where is information located?, is quite different. The primary reason for this is because Information-A conceptualisations of information emphasise the semantic nature of information. As a consequence, Information-A conceptualisations of information take information to be 'about' something, whether that something is possible ways the world could be (information as range), a given situation in the world (systemic approach), or the true ways that the world is (veridical approach). This is the case because all Information-A conceptualisation of information presuppose that information only makes sense within the context of the meaningful activities of an agent. That is, that information must be conceptualised as something used by agents in relation to questions of error, truth, and knowledge. Where information is to be located for Information-A conceptualisation, then, is only of secondary importance to the question, how does information connect to meaning?

Because one's conceptualisation information must derive from one's meaningful engagement with modal space, a given situation, or truth, and so it only makes sense to ask where information is to be located from within the context of this meaningful engagement itself. Thus, on the Information-A account, information must be 'about' something before it can be located 'in' something: information is first and foremost a semantic phenomena.

Now, this is not to say that Information-A conceptualisations of information cannot ask the question, where is information located?, after the establishment of meaningful activities from which to derive their conceptualisation of information. But seeing as this question emerges from the semantic conceptualisation of information it can, in fact, be reduced to the second of the questions I posed above; namely: Can there be information without an informee? Or, put differently: Can there be information without the semantic activities of agents? As Floridi (2011) put it:

> The problem here is whether there is information in the world independently of forms of life capable to extract it and, if so, what kind of information is in question.

The question, can there be information without an informee?, is an open question for Information-A, Information-B, and Information-C conceptualisations of information. To see why, first consider the position that holds that there is information without an informee: environmental information. Environmental information is the position that holds that information is in the world (they say 'environment') independent of the 'representational' and semantic activities of higher level organisms.[68] Dretske (1981), for example, argues that environment information is independent of the observer (informational agent) even if it is defined relative to an observer (informational agent). The standard example given in the literature is of the series of concentric rings visible in the wood of a cut tree trunk, which can be used to estimate the trees age. For Dretske, the information about the age of the tree is 'in' the world regardless of whether there are any beings capable of higher level cognitive – i.e. 'representational' or semantic – functions. Moreover, as Floridi again points out, aspects of the world such as:

> Plants (e.g., a sunflower), animals (e.g., an amoeba) and mechanisms (e.g., a photocell) are certainly capable of making practical use of environmental information even in the absence of any (semantic processing of) *meaningful* data (Floridi, 2011)

Still, the definition of information is always given relative to an observer (informational agent).[69] In other words, the information contained 'in' the series of concentric rings visible in the wood of a cut tree trunk is only taken to be about the tree's age when defined relative to the semantic

---

[68] The term representation here is used in the psychological sense of being an internal cognitive symbol – or mental state – that represents external reality. Even if we take the mental state to be the representation-device, this would still not square with my [REP] conceptualisation of representation in chapter 1 however (although perhaps it could be made compatible given an interpretation of the mental state as a representation-device as just mentioned). I use the term 'representation' here, then, merely to embed myself into the literature and not to be consistent with what has preceded this chapter in chapters 1 and 2.

[69] It should be said that to be able to give a *definition or conceptualisation of information*, observers (information agents) must be capable of higher level cognition; i.e. representation and semantics. But this does not necessarily imply that all information agents need to possess higher level cognition – the plant in Floridi's example certainly does not, for instance.

framework of an observer who can meaningfully talk 'about' age, trees, concentric rings etc. And this is where we reach our problem in regards to the different conceptualisations of information I introduced above.

The problem is that Information-A, B, and C conceptualisations of information do not specify whether they take information to require an informee. As I have said, Information-B and Information-C conceptualisations of information take information to be 'in' messages or strings respectively. But this in itself says nothing about whether the existence of messages and strings themselves presupposes the existence of observers (information agents) capable of higher level cognition. Similarly, Information-A conceptualisations of information take information to always 'about' something, where being about something is connected to the semantic activities of observers (information agents) capable of higher level cognition. But, again, this claim in itself says nothing about whether information can exist independent of being about something, since it may only be the case that an Information-A *definition* or *conceptualisation* of information is 'about' something relative to the semantic framework of the observer (information agent).[70] So neither Information-A, B, nor C in and of themselves explicitly stipulate whether or not the *existence* – or ontology – of information is relative to the existence of observers (information agents) capable of higher level cognition.

I have said that the difference in Information-B/C and Information-A responses to the question, where is information located?, are such that one side says that information is 'in' something (Information-

[70] Above we saw that some Information-A conceptualisations of information do hold that information is prior to cognitive action (e.g. Barwise and Perry, 1983).

B and C, messages and strings respectively) and the other side says that information is first 'about' somethings (Information-A). But neither Information-A, B, or C necessarily gives a conceptualisation of information that specifically takes information to exist independent of observers (information agents). Such a conceptualisation of information as self-subsistent is possible (environmental information), and so the fact that neither Information-A, B, nor C can be said to explicitly conceptualise information in this way is, I think, revealing. To get clear on what is revealed exactly, I will now give my own interpretation of the *pre-theoretical bases* of Information-A, B, and C conceptualisations of information.

Pre-theoretical Bases: Substantive/Extensionalist vs. Semantic/Intentionalist?

I have said that neither Information-A, B, nor C concepts of information *explicitly commit* to any stance in regards to the question, can there be information without an informee? But that is not to say that philosophers have remained silent about the matter. We have already seen, for instance, how advocates of environmental information like Dretske and Barwise do think that information can – and does – exist without an informee. And philosophers like Fodor (1987) and Searle (1980, 1990) are strong advocates of the view that information depends, in some form, on intelligence. So what is the central difference between these two camps?

The answer, I think, is that one camp argues that information can be naturalised (Dretske, Barwise), whereas the other does not (Fodor, Searle). What it means to say that information can be 'naturalised', however, remains unclear. For my purposes, then, I will say that the difference between the likes of Dretske and Barwise on the one hand, and the likes of Fodor and Searle on the other, is a difference shaped by two complementary dichotomies: substantive vs. semantic conceptualisations of information and extensionalist vs. intentionalist conceptualisations of information. The substantive vs. semantic conceptualisation of information is a dichotomy between conceptualising information as a *thing* – a substance – 'in' reality, or conceptualising information as a property of semantic activities that are 'about' the world in some sense; i.e. meaningful discourse or science. The extensionalist vs. intentionalist conceptualisation of information is a dichotomy between conceptualising information as something to be picked out in the world as the *ex*tension of the concept "information", or conceptualising information as something dependent upon the mind – and hence *in*tentions – of the agent. Clearly, then, the two dichotomies can be merged into one; i.e. substantive/extensionalist vs. semantic/intentionalist.

To make this broader dichotomy clear, consider the example of the Rosetta Stone which made possible our understanding of Egyptian hieroglyphics. On the substantive/extensionalist account, Egyptian hieroglyphics were information for eighteenth century man even prior to the discovery of the Rosetta Stone, because information is conceptualised as something 'in' the world and Egyptian hieroglyphics are one of the many extensions of the concept information that could, in principle, be picked out. But on the semantic/intentionalist account, Egyptian hieroglyphics were *not* information for eighteenth century man prior to the discovery of the Rosetta Stone, because information must always be 'about' the world and hence is dependent on the semantic and intentional activities of man, and this 'aboutness' was not possible in the case of the

Egyptian hieroglyphics for eighteenth century man prior to the discovery of the Rosetta Stone.

Now, the debate between the substantive/extensionalist conceptualisations of information and the semantic/intentionalist conceptualisations of information occupies its own place in the philosophical discussion of information as a whole. But I think this discussion has influenced the development of the various Information-A, B, and C conceptualisations of information. The shape of this influence, I think, is such that the various Information-A, B, and C conceptualisations of information have been developed with the *values* of either the substantive/extensionalist conceptualisation of information or the *values* of the semantic/intentionalist conceptualisation of information in mind. That is, that the various Information-A, B, and C conceptualisations of information are manifestations of *implicit* commitments to a view of information that squares *either* with the substantive/extensionalist conceptualisation *or* with the semantic/intentionalist conceptualisation of information.

Take Information-B and Information-C conceptualisations of information to begin with. Both agree that information should be located 'in' something (messages or strings). And why do they think this? One reason is certainly because it enables a quantification of information in terms of probability space. But this quantification can only get going once information has been conveniently *given substance*. By 'given substance' here I do not necessarily mean physicalized. But I do mean to imply that information is conceptualised as something that is able to be 'in' something else (a message or string, for instance) and as something that is composed of something else (data). This, of course, is the typical way of thinking about ordinary objects given a rudimentary interpretation of our current scientific image.[71] And so by thinking of information in these terms as well, one conceptualises information as something consistent with our typical scientific-image based conceptualisation of the world. What's more, perhaps the most important feature of our contemporary scientific-image based conceptualisation of the world is the commitment to the idea that a concept should have extensions 'in' the world to be permissible (it should be empirically verifiable). And the substantial conceptualisation of information offered by Information-B and Information-C adheres to this commitment, because the concept of information is taken to have an extension in the (data-)content of any identifiable message or string.

Following from what has just been said, therefore, I submit that the reason that Information-B and Information-C conceptualisations of information developed as they did owes in a large part to an implicit commitment to substantive/extensionalist conceptual commitments. I do not want to give any arguments for why such an implicit commitment would obtain, but I only want to state clearly that I do think that such a commitment does obtain.[72] In this way, the conceptualisation of information offered by Information-B and Information-C is buttressed by

---

[71] I should state clearly that I think this kind of interpretation of our current scientific image reflects a misunderstanding of the results of contemporary science. But that does not stop many philosophers from thinking in such a manner, as Ladyman, Ross et al. have identified to humorous effect with their quip about the 'philosophy of A-level chemistry' (c.f. Ladyman, Ross, et al., 2007: 24).

[72] The reasons for such a commitment will likely be a combination of naturalistic, physicalistic, scientistic, and other factors. But I leave the reader to decide for him/herself whether such factors would or would not be appealing, and whether such factors are or are not valid.

substantive/extensionalist conceptual commitments *that are not to be found explicitly in Information-B or Information-C conceptualisations of information themselves*. In other words, the conceptualisation of information offered by Information-B and Information-C are *based* on pre-theoretical commitments to the substantive/extensionalist conceptualisation of information, without explicitly demonstrating that they are substantive/extensionalist in conceptual nature.

Turning our attention now to Information-A conceptions of information, we can recall that information was conceptualised as meaningful and hence 'about' the world in some sense. But, again, we must ask why Information-A typically conceptualises information in this way. It is quite evident that in order to conceptualise something as being 'about' something else one first requires some appreciation of what it means to have observers (information agents) capable of higher level cognition. Moreover, if information is to be meaningful, then one also requires some appreciation of the semantic capabilities of observers (information agents). It follows, then, that any conceptualisation of information as 'about' something will always issue from an understanding of the way that the higher level cognition of observers (information agents) functions. But just how the higher level cognition of observers (information agents) actually functions remains a difficult and divisive question in contemporary science, and so it is not surprising that there is no consensus about what it means to say that Information-A conceptualisations of information issue from an understanding of cognitive and semantic functioning of agents.

One option here is to argue that information should be conceptualised just as a part of the higher level cognition and semantics of observers (information agent). One could read the information as range conceptualisation of information in this way, because information is conceptualised as the range of options *an agent* has for how the world could be. Similarly, one could read the veridical approach to semantic information in much the same way, because information is taken to be truthful and truth, arguably, is inseparably bound to the higher level cognition and semantics of observers (information agents). But one may also be critical of the idea that information should be conceptualised just as a part of the higher level cognition and semantics of observers (information agents). One could argue, for instance, that information is independent of the higher level cognition and semantics of observers (informational agents) even if it is conceptualised relative to the higher level cognition and semantics observers (informational agents). This is exactly the position most systemic approaches to information – like environmental information – endorse. On the systemic view, therefore, information is actually 'in' the environment, but any conceptualisation of information must issue from the way that higher level cognition of observers (informational agents) is 'about' the environment.

Although there are differences between the different versions of the Information-A conception of information, I still hold that *all* have been influenced by the semantic/intentionalist conceptualisation of information. I think that this influence has been such that all Information-A conceptualisations of information *implicitly* endorse the claim that any viable conceptualisation of information must at least issue from the higher level cognitive functions of observers (information agents). As we have seen, however, the options are then open as to where information is located – either in the cognitive function itself (the mind) or in the environment (the world). But what is not an option for Information-A conceptions is to go about conceptualising information in such a way that discounts semantic and intentional factors altogether. In this way, Information-A conceptions of information subscribe to an implicit commitment as to how information should be conceptualised; namely by first paying attention to the higher level cognitive functions of observers (information agents) and then developing further conceptualisation of what information *is*. Again, this is not to say that all Information-A conceptions of information will agree that information is necessarily bound to the semantic and intentional activities of agents. But it is to say that that Information-A conceptions of information are based on pre-theoretical commitments to the semantic/intentionalist conceptualisation of information *to the extent that the conceptualisation of information they give always stems from the semantic and intentional activities of agents*.

It is this distinction between substantive/extensionalist and semantic/intentionalist pre-theoretical bases, I argue, that is the most fundamental divide between different conceptualisations of information. I think that this is the case because at the highest level of generality we are able to group together *all* Information-B and C conceptualisations of information by reference to their commitment to a substantive/extensionalist pre-theoretical base, and we are able to group together *all* Information-A conceptualisations of information by reference to their commitment to a semantic/intentionalist pre-theoretical base. But the dichotomy between the substantive/extensionalist and the semantic/intentionalist pre-theoretical bases prevents a grouping together

– or reconciliation – of *all* Information-A, B, and C conceptualisations of information.

Beyond the Dichotomy?

The conceptualisation of environmental information by Dretske (1981) and the conceptualisation of information as semantic content by Bar-Hillel and Carnap (1953) embody the best attempts to move beyond the dichotomy between the substantive/extensionalist and the semantic/intentionalist conceptualisations of information.

We have already seen how Dretske (1981) advocates a position where information is 'in' the environment, but this conceptualisation of information is taken to issue from the way that observers (informational agents) think 'about' the environment. The example that makes this clearest is the starry sky in which the same pattern of stars, at different moments in time and locations in space, will in general convey different information about one's own or another person's location in time and space. For Dretske, information is 'in' the starry sky itself, but our appreciation of this information is always relative to some features of our own cognition; i.e. our semantic framework. Thus, Dretske conceptualises information in the substantive/extensionalist sense, but is aware that to reach this point one must first consider the mediation of one's own higher level cognitive function – a typical semantic/intentionalist idea. So the argument is that a substantive/externalist conceptualisations of information can only be given as a consequence of semantic/intentionalist factors.

Bar-Hillel and Carnap take the opposite route to Dretske by beginning with a conceptualisation of information that is strongly substantive/extensionalist, but then transforming this position until it becomes more akin to a semantic/intentionalist conceptualisation.[73] This approach has been termed the 'probabilistic approach to semantic information', because of both its affinity with Shannon's conceptualisation of information and its ultimately semantic nature (c.f. Adriaans and van Benthem, 2008). Bar-Hillel and Carnap argue that 'semantic content (CONT) in [a proposition] p is measured as the complement of the *a priori* probability [(Pr)] of p' according to this formula (Floridi, 2011):

$$\mathrm{CONT}(p) \;=\; 1 \;-\; Pr(p)$$

Where the informativeness (IN) of p is calculated, following Shannon's equation, $H(Pr) = -\Sigma_{(x\,\in\,A)} Pr(x) \log_2 Pr(x)$, such that for the equation just given above, the reciprocal of $\mathrm{Pr(p)}$, expressed in bits, where $\mathrm{CONT}(p) = 1 - Pr(p)$, is:

$$\mathrm{IN}(p) = \log \frac{1}{1 - \mathrm{CONT}\ (p)} = -\log Pr(p)$$

Now, the fact that informativeness is given by a probability measurement expressed in bits seems to indicate that Bar-Hillel and Carnap are engaged in a substantive/extensionalist conceptualisation of information. But things are complicated by the fact that, for Bar-Hillel and Carnap, 'the probability distribution is the outcome of a logical construction of atomic statements according to a chosen formal language' (Ibid.) And so although

[73] For further developments of Bar-Hillel and Carnap's work see Kemeny, 1953; Smokler, 1966; Hintikka and Suppes, 1970; and even Dretske, 1981 (of course Dretske's work shifts the attention to the environment or situation).

information is conceptualised via a probability space expressed in bits, it is still the case that this probability space itself is relative to the higher level cognitive functioning of observers (information agents) who are able to construct a logico-probabilistic space around certain *a priori* atomic statements. Information, then, is conceptualised by Bar-Hillel and Carnap in probabilistic terms in much the same way as Information-B conceptions of information. But, unlike other Information-B and Information-C conceptualisations, Bar-Hillel and Carnap suppose that the very possibility of undertaking such a probabilistic conceptualisation of information requires first that we have a semantics for a formal language that facilitates the logical construction of a probability space based on atomic statements.[74]

Thinking again of the Rosetta Stone example above, we can elucidate both Dretske's and Bar-Hillel and Carnap's conceptualisation of information. For Dretske, information is certainly 'in' the Egyptian hieroglyphics prior to eighteenth century man's acquisition of the Rosetta Stone (environmental information). However, it is only with the acquisition of the Rosetta Stone that Dretske would say that eighteenth century man was able to access, define, and conceptualise the environmental information scattered all across Egypt, because only then was eighteenth century man able to align the information in their environments with the semantics of their higher level cognition. Still, information itself is, for Dretske, fundamentally substantive/externalist.

---

[74] The basic point here is that Carnap's probabilistic conception of semantic information relies upon an *a priori* measure which is extremely language dependent; i.e. dependent on the atomic statements that make possible a formal language.

For Bar-Hillel and Carnap, information is also taken to be 'in' the Egyptian hieroglyphics, but only insofar as such Egyptian hieroglyphics are able to be assigned a probability in accordance with the semantics an agent has for a given formal language. In other words, any p that is a particular sequence of hieroglyphics will be informative if p can be assigned a particular probability, where the assigning of a probability to p is relative to parameters that are set by a logical framework constructed by *a priori* atomic statements that brings with it a certain semantics that attaches to the higher level cognition of observers (information agents). So probability is derived from deductive reasoning that is valid given a semantics for the logical framework in question. Thus, although the probability appears to be 'in' p, it is in fact only 'in' p to the extent that p can be assigned a probability, and this first requires that the observation language of p is made to correspond to the formal language that gives the semantics for the logical framework in question. So information can only be taken to be 'in' p so long as the product of the higher level cognition of the observer (information agent) – that is, the formal language and logical framework – can be said to be 'about' p. Put differently, information can only be taken to be 'in' the Egyptian hieroglyphics so long as the semantics of the formal language being used to calculate the information 'in' the hieroglyphics corresponds to, or is known to be meaningfully 'about', the hieroglyphics. And so, for Bar-Hillel and Carnap, eighteenth century man would only have been able to say that information was 'in' the hieroglyphics after the acquisition of the Rosetta Stone that enabled a correspondence between eighteenth century man's observation and formal languages, and hence the assigning of probability to any particular sequence of hieroglyphics p. Ultimately, therefore, Bar-Hillel and Carnap still advance a semantic/intentionalist conception of information itself.

Here, then, we have two cases that attempt to mesh the substantive/extensionalist and the semantic/intentionalist conceptualisations of information, and ultimately come up short. These cases are, however, at least attempts to break free of the pre-theoretical commitments that hold apart substantive/extensionalist and the semantic/intentionalist conceptualisations of information. That is, they are at least attempts to give conceptualisations that account for the interrelation of the way that information is both 'in' the world and 'about' the world at the same time. Despite the focus on the interrelation of 'in' and 'about' the world, however, both Dretske and Bar-Hillel and Carnap ultimately come down on one side or another of the substantive/extensionalist and the semantic/intentionalist divide. What is perhaps most disconcerting is that these two more nuanced conceptualisations of information offered by Dretske and Bar-Hillel and Carnap exemplify the best philosophical attempts to break free of the pre-theoretical commitments that hold apart substantive/extensionalist and the semantic/intentionalist conceptualisations of information. Aside from these two positions – and some other modifications of these two positions (c.f. footnote 59 above) – I know of no other philosophical conceptualisations that even attempts to move beyond the substantive/extensionalist and the semantic/intentionalist dichotomy.

Evidently, then, formulating such a conceptualisation of information as both 'in' the world and 'about' the world at the same time is problematic. As the most cutting edge engagement with information illustrates (e.g. Floridi, 2011), philosophers find it difficult to think of information as being 'in' and 'about' the world without marking a difference between something 'in' the world and the way that the cognitive processes of information agent are 'about' the world.[75] This, I think, is the central problem that any novel conceptualisation of information must grapple with and as yet remains unresolved.

What Information Must Be

Before moving on to offer my resolution to the problematic dichotomy between the substantive/extensionalist and the semantic/intentionalist conceptualisations of information, it is helpful to consider the requirements that philosophers now agree any conceptualisation of information must adhere to. These requirements are:

> that information ought to be quantifiable (at least in terms of partial ordering), additive, storable and transmittable. (Floridi, 2008: 118)

The first of these requirements – that of being quantifiable – is satisfied by conceptualisations of information as both 'in' and 'about' the world; i.e. Information-A, B, and C. This is the case because it is possible to give quantity to information taken as both 'in' and 'about' the world by making use of either the relevant mathematics related to probability space; the

---

[75] In particular, I am referring here to Floridi's distinction between data as '*differentiae de re*, that is, mind-independent, concrete points of lack of uniformity in the fabric of Being', and information as *differentiae de re* that are then 'epistemically exploit[ed] as resources, by agents like us, for their cognitive processes' (Floridi, 2011: 368). For all intents and purposes, Floridi is marking a difference between what is 'in' the world independent of observers (i.e. data) and what comes about with observers cognitive activities 'about' the world (i.e. information).

relevant mathematics related to computation; or the relevant logic related to doxastic (i.e. belief related) modal spaces of one kind or another (epistemic, intensional). And with the quantifiable requirement satisfied it becomes possible to make sense of the additivity, storability, and transmissibility of information *relative to the way that information is quantified.* That is to say, that Information-B and Information-C conceptions of information (i.e. Shannon's communicative theory of information and Kolmogorov complexity) will account for the additivity, storability, and transmissibility of information in terms of the relevant mathematics related to probability space. Whereas Information-A conceptions of information will account for the additivity, storability, and transmissibility of information in terms of the relevant logics related to doxastic (i.e. belief related) modal spaces of one kind or another.

Take the case of both storability and transmissibility. It may seem *prima facie* permissible to say that all conceptions of information take information to be stored in messages, strings, or propositions, and so take information to be transmitted following from the communication or computations of these messages, strings, or propositions. Such an account of the storability and transmissibility of information, however, tells only half the story, because it does not make explicit the fact that messages, strings, and propositions can only be identified relative to a particular quantification of information. Consider an instance of information σ, where σ is stored in some message, string, or proposition τ and is able to be transmitted given the communication or computation of the message, string, or proposition τ. How, we must ask, are we first able to apprehend or come into contact with τ as something informative? One could answer that we simply perceive that τ is informative by means of one or another of our sense modalities. But I do not think this answer is sufficient. A whole host of people perceived the Egyptian hieroglyphics without finding them to be informative, for example, and so it seems that something more than mere perception is required for us to identify any message, string, or proposition τ as an instance of information σ.

Some potential messages, strings, or propositions τ, then, will not be recognisably informative given our sense modalities alone – perhaps because the message, string, or proposition τ seems disordered (random even), nonsensical, occurs very infrequently, or is incapable of being materially realised (abstract). In cases like these the very possibility of taking a message, string, or proposition τ to be informative depends upon the method that one has available to measure the information contained 'in' the message, string, or proposition τ or the information that the message, string, or proposition τ is 'about'. For Information-B conceptualisations of information, this requires the mathematics related to probability space.[76] For Information-C conceptualisations of information, this requires the mathematics related to computation. And

[76] Interestingly, another potential shortcoming of Shannon's communicative theory of information derives from the consequence of a probabilistic measure of specifying informative messages. The reason for this is that even if one is able to specify informative messages, one may be unable to differentiate a collection of specified informative messages based upon their content, because this content can only be measured in bits. Take the messages 'CAT', 'JOB', and 'XGK', for example. All three may be specified as informative given a probabilistic measure that confines itself to the letters of the alphabet (A-Z), and so is concerned with 27 symbols (all the letters and an additional symbol for a space). But by making use of the negative base-2 logarithm to work out the probability of any single symbols occurrence, we find that there is 4.76 bits of information corresponding to 4.76 bits of uncertainty or 4.76 bits of Shannon entropy per symbol ($I = -\log_2 {}^{1}/_{27} = 4.76$ bits per symbol). And, thus, that 'CAT', 'JOB', and 'XGK' all store and are capable of transmitting the same amount of information (measured in bits) ($I = 3 \times (-\log_2 {}^{1}/_{27} = 14.28$ bits of information)).

for Information-A conceptualisations of information, this requires logic related to, for instance, doxastic (i.e. belief related) modal spaces of one kind or another. What this implies, therefore, is that one's identification of what counts as an informative message, string, or proposition is to at least to some degree *quantification relative*; i.e. an informative message, string, or proposition can only be specified as being informative given a particular quantitative measure of what it means to be informative. In this way, the messages, strings, or propositions that one takes to be the transmissible storage 'containers' or 'bearers' of information also depends on a particular measure of what it means to be informative – and this measure, according to the requirement set by philosophers at least, should be quantitative.

Additivity too is quantification relative. The reason, much like the cases of storage and transmissibility, is that to add separate instances of information σ and σ_2 together one must first have measured both σ and σ_2 in such a way that the operation, $\sigma + \sigma_2$, is even possible. In other words, one must first have given commensurable measurements of both σ and σ_2, and, of course, to give such compatible measurement requires that one first adopt a particular quantitative measure of what it means for any message, string, or proposition to be the container or transmitter of an instances of information σ and σ_2.

The requirements of additivity, storability, and transmissibility, therefore, can be satisfied on both conceptions of information 'in' the world and conceptions of information 'about' the world. The reason for this is because they are quantification relative and all Information-A, B, and C conceptualisations of information support a quantification of information in one way or another. So the requirements themselves – however helpful in setting out what any conceptualisation of information must adhere to – provide no assistance in helping us to move beyond the problematic dichotomy between the substantive/extensionalist and the semantic/intentionalist conceptualisations of information.

In fact, I think the requirements currently under discussion may have inadvertently had a negative impact on the development of novel conceptualisations of information that *do* move beyond the dichotomy I have been discussing in this chapter. The reason for this is quite simple: I think the requirements may have had a negative impact on the development of novel conceptualisations of information that move beyond the dichotomy I have been discussing, because any philosopher who gives a conceptualisation that satisfies the requirements may be inclined to think his conceptualisation is as viable as any other conceptualisation available. That is, any philosopher whose conceptualisation of information meets the requirements of quantifiability, additivity, storability, and transmissibility may feel that she has *succeeded* in conceptualising information.

This state of affairs, I submit, is a mass delusion perpetuated by most – if not all – contemporary philosophers of information, because even if a conceptualisation of information can be said to satisfy the requirements of quantifiability, additivity, storability, and transmissibility, it will probably still come down on one side of the dichotomy between substantive/extensionalist and semantic/intentionalist conceptualisations of information. Such 'successful' conceptualisation of information, then, still fail to address the fundamental dichotomy in regards to our understanding of information as a phenomenon and a concept: the question of whether information is 'in' the world, 'about' the world, or somehow both. I hold, therefore, that whilst any conceptualisation of information must conceptualise information as quantifiable, additive, storable, and transmissible, it must first and foremost conceptualise information in such a way that resolves, invalidates, or in some way moves beyond the 'in' the world/'about' the world dichotomy I have been discussing. Only then, I think, will we begin to get a clearer picture of what information *is*.

Conceptualising Information as a Structural Relationship

In this chapter, I first gave an overview of different conceptualisations of information, before illustrating that these conceptualisation are divided between commitments to either substantive/extensionalist or semantic/intentionalist conceptualisations of information. Then, I considered attempts to move beyond the substantive/extensionalist and semantic/intentionalist divide, and found that up to this point no philosophical conceptualisation of information has succeeded in overcoming the substantive/extensionalist and semantic/intentionalist dichotomy. Following that, I introduced the standard requirements that any conceptualisation of information must satisfy, but noted that such requirements do not by themselves resolve the substantive/extensionalist and semantic/intentionalist dichotomy.

Given this state of affairs, I will now develop a novel conceptualisation of information that moves beyond the dichotomy between the 'in' the world – substantive/extensionalist – and 'about' the world – semantic/intentionalist – conceptualisations of information. Moreover, I will explain why the conceptualisation of information I endorse can also satisfy the requirements of quantifiability, additivity, storability, and transmissibility.

The conceptualisation of information I endorse is as follows:

[INF] Information is the objective structural relationship between (the semantic activities of) intentional agents, the information-bearer, and the world.[77]

[INF], then, is a conceptualisation of information *as* a structural relationship between semantic agents who ask questions such as "How can something count as information? And why?", "how is information related to error, truth and knowledge?", and, "when is information useful?"; an information-bearer, which is any message, string, or proposition formulated according to rules (syntax) that govern the system, code or language in question[78]; and the world.[79] As we have seen above, however, we can capture the relation between agents and world by simply

---

[77] I use the word 'structural' once again to express the idea that all three components are interrelated and support one another as one organisationally interconnected whole: to reiterate, there is *one* structural relationship with three components.

[78] Information-bearers, it is worth reiterating, need not be linguistic entities. A painting or a model, for example, is also formulated according to some rules and so will also count as an information-bearer.

[79] Exactly what is meant by 'the world' in [INF] (and [REP], if you recall the end of chapter 1) is left open and will not be dealt with until the next and final chapter. I will make one point now, however. I am committed to the claim that there is a world 'outside' of the mind, and so I reject the radical idealist position that there is no world independent of the semantic activities of agent (i.e. Berkeley, 1948-1957). The question of the importance of this mind-independent world to representation and information, however, remains open for debate (see my conclusion below).

referring to logical space – jointly constituted and instantiated by agent and world – that agents make use of to think about the world ([LOGSP*]),[80] and hence we have [INF*]:

[INF*] Information is the objective structural relationship between a logical space ([LOGSP*]) and the information-bearer.

The first important point to note about [INF*] is that it conceptualises information as both 'in' and 'about' the world at the same time. Information, according to [INF*], is 'about' the world, because the structural relationship that *is* information is only possible when there are intentional agents capable of the kind of higher level cognitive activities that bring with them meaningful practices related to, amongst other things, belief and truth. But information is also 'in' the world to the extent that [INF*] suggests that the structural relationship that *is* information is only possible when there is a world, alongside intentional agents, to jointly constitute and instantiate logical space ([LOGSP*]).

Now, let's return once again to the Rosetta Stone example to elucidate how [INF*] conceptualises information. On my [INF*] conceptualisation of information, information can neither be said to be 'in' the Egyptian hieroglyphics nor 'about' the Egyptian hieroglyphics separately. Thus, prior to the discovery of the Rosetta stone, there was simply no information 'in' or 'about' the hieroglyphics at all, because there was no structural relationship between the information-bearer (chiselled into the walls of pyramid, carved into pieces of wood, or written on papyrus) and a logical space ([LOGSP*]) jointly constituted by agent

[80] See page 88 above.

(eighteenth century man) and world (c.f. footnote 64 above). Whilst it certainly was the case that at one time there was such a structural relationship between information-bearer and logical space – i.e. *circa* 3500 BCE–400 CE, during the eighteenth century this structural relationship, and hence the information in question, simply did not exist. Thus, the Rosetta Stone should in fact be seen as one of the things that brought into being certain information about Ancient Egypt *at the end of eighteenth century*. According to [INF*], therefore, such information – whether defined in terms of bits of data (Information-B/C) or meaningful messages (Information-A) – was only brought into being because the Rosetta Stone made possible the obtaining of a structural relationship between logical space ([LOGSP*]) and the information-bearers.

But now one could ask: How was the informational structural relationship ([INF*]) between logical space ([LOGSP*]) and information-bearer (Rosetta Stone) brought into being? In other words, how was the structural relationship that brought with it information about Egyptian priests, the Macedonian Pharaoh Ptolemy V Epiphanes, and the siege of Lycopolis that is carved into the Rosetta Stone itself brought into being? The answer one gives will inevitably have to refer to other objects such as Greek and Demotic texts involved in the deciphering of the Rosetta Stone etc. But then we seem to require an explanation about how Greek and Demotic texts brought information about Ancient Greece and Ptolemaic Period into being, and so on *ad infinitum*. Thus, the stage is set for a *reductio* argument against the [INF*] conceptualisation of information as a structural relationship. For my part, I completely accept this *reductio* argument, but think that it is anything but vicious. In fact, I think that ultimately the *reductio* in question merely amounts to asking the question:

Where did information come from?, which is equivalent to asking a semanticist a question such as: Where did meaning come from?; or asking a psychologist a question such as: Where did the mind come from? For this reason, I do not take the *reductio* argument under consideration to be one worth worrying about: the structural relationship between logical space ([LOGSP*]) and the information-bearer that *is* information just is part of reality as we find it.

Returning to the point, we have seen that information as conceptualised by [INF*] is brought into being at least in part by artefacts and objects such as the Rosetta Stone. The reason for this is clear: because objects and artefacts like the Rosetta Stone are information-bearers. But we must be careful not to forget that according to [INF*] information can only be brought into being as a structural relationship *between* logical space ([LOGSP*]) and information-bearer. To be clear, then, the Rosetta Stone as a chunk of rock could never be enough to bring into being information about Ancient Egypt alone, because as an information-bearer the Rosetta Stone is a necessary but not sufficient condition of information. If, for example, a meteor had hit earth in the mid seventeenth century and annihilated all observers capable of higher level cognition (information agents), then the Rosetta Stone would certainly not have been able to bring information about Ancient Egypt at the end of the eighteenth century into being. To put what I am gesturing at simply here, informational structural relationships require as a sufficient and necessary condition that agent, information-bearer, and world all exist, because without all three we cannot have the structural relationship *between* logical space and information-bearer. For [INF*], then, information certainly requires an informee, if by informee we mean intentional and semantic agent. But information is not be located in the semantic activities of agents alone, because information is a structural relationship *between* information-bearer, agent and world (i.e. logical space ([LOGSP*])).

Having said, then; (a) that informational structural relationships require as a sufficient condition that agent, information-bearer, and world all exist; and (b) that information *is* the objective structural relationship between information-bearer and logical space ([LOGSP*]); we can draw the most important conclusion about [INF*]. This conclusion is as follows: according to [INF*], information is that which is enacted when information-bearers enter into a structural relationship with agent and world (i.e. [LOGSP*]). This *enacting*, it is vital to note, is no more the result of information-bearer, world, or agent taken individually. Instead, the combination of agent and world (logical space) and information-bearer come together to *enact* the structural relationship that *is* information. Thus, [INF*] is the only conception of information currently available that moves beyond the substantive/extensionalist and semantic/intentionalist divide, by reconciling the substantive/extensionalist and semantic/ intentionalist. For this anti-dualistic reason, I think it should be vigorously defended.

The final problem, however, is whether or not [INF*] can satisfy the requirements of quantifiability, additivity, storability, and transmissibility as other conceptualisations of information all do. I propose that [INF*] can satisfy these requirements by reinterpreting practically any form of Information-A conception of information in accordance with the [LOGSP*] account of logical space I gave above. To remind ourselves once again, I defined logical space on page 93 above as:

[LOGSP*] Logical space is jointly constituted and jointly instantiated by agent(s) and world.

And so it follows that if we take [INF*] to endorse a quantification of information in terms of a logical space jointly constituted and jointly instantiated by agent(s) and world ([LOGSP*]) – such as, for instance, doxastic (i.e. belief related) modal space –, then [INF*] can be said to also satisfy the requirements of additivity, storability, and transmissibility that follows from the Information-A quantification of information in terms of a logical space. [INF*], then, can satisfy the requirements of quantifiability, additivity, storability, and transmissibility, but only if [INF*] is taken to endorse a logico-semantic quantification of information in terms of a logic that codifies the logical space [LOGSP*] of possibility (modal logic), belief (doxastic space), obligation and permission (deontic logic), etc.

In taking [INF*] to endorse a logico-semantic quantification of information, however, we must be careful specify what [INF*] takes a logic to be. For one thing, [INF*] must take a logic to be the means of quantifying information if it is to count as a version of Information-A conceptions of information at all. But the [INF*] conceptualisation of information will only be possible if logic is also taken to be an information-bearer that carries or communicates information.[81] To see how this is

[81] I think it is fair to say that the information-bearer is also part of the world (c.f. footnote 2 above). But if one does not agree with this claim then where, I ask, are we to locate the information-bearer exactly? 'In' the agent's semantic activities? If this is the position you are inclined to adopt, then it is still the case that a reinterpretation of logic as both a codification of logical space and as an information-bearer is compatible with [INF], because both information ([INF]) and logical space ([LOGSP*]) are taken to be jointly constituted and instantiated by the semantic activities of the agent and the world. And if one wants to say that the information-bearer is neither part of the world nor the agent's semantic activities, then the burden of proof is no longer on me,

possible we must realise that from the perspective of [INF*] information is *made possible* by logic as a semantically imbued information-bearer (as well as other semantically imbued information-bearers like the concentric rings of a tree or ordinary linguistic content). But the logic and other information-bearer's are *not* equivalent to information, because such information-bearers are merely natural objects and/or artefacts that enter into a structural relationship with logical space. In this way, logic is a special kind of information-bearer for the [INF*] conceptualisation of information, because it does two things at once. Firstly, it makes possible the quantifiability, additivity, storability, and transmissibility of information. And, secondly, it makes possible information by standing in a structural relationship to logical space in the [LOGSP*] sense. [82] In this respect, according to [INF*], logic – like the models of science in the last chapter and other information-bearing objects and artefacts like newspapers and concentric rings in tree trunk – enters into the structural relationship with logical space ([LOGSP*]) that *is* information. This is not to say that the structural relationship between information-bearer and logical space ([LOGSP*]) did not exist prior to the development of logic (it certainly did), but only that logic is a privileged kind of information bearer that allows for us to comprehend how information as a structural relationship is quantifiable as well as being additive, storable, and transmissible.

---

because one will have to work hard to explain exactly where and what the information-bearer *is* on this picture.

[82] I realise, of course, that to accept this proposal requires that philosophers adopt an approach to logic, logical space, and information that leaves the ghost of Cartesian dualism well behind. Many, I am sure, will have trouble shaking the ghost even where they claim to be working in an anti-dualist tradition.

Chapter 4.

Informational Scientific Representation

In this final chapter I will bring together the three concepts I have discussed in the preceding three chapters; namely, representation, scientific representation, and information. First, I will defend a conceptual synthesis of the concepts of representation and information as I have defined them in this essay. I call the result of this conceptual synthesis: informational representation. Then, I explain how a key feature of informational representations is the presenting of logical space by an object or artefact, followed by the making sense of how the object or artefact presents logical space. This leads to consideration of the link between the making sense, accuracy, and perspective. And, ultimately, to my introduction of one kind of informational representation in particular: informational scientific representations. Here, I argue that the structural relationship that *is* informational scientific representation differs from other kinds of structural relationships because it is guided by the aim to give the most accurate approximation of reality possible.

Conceptual Synthesis

We can begin by first recalling my conceptualisation of representation in chapter 1 and information in chapter 3 in turn:

[REP] Representation is the objective structural relationship between (the semantic activities of) intentional agents, the representation-device, and the thing represented.

or

[REP*] Representation is the objective structural relationship between logical spaces ([LOGSP*]) and the representation-device.

and

[INF] Information is the objective structural relationship between (the semantic activities of) intentional agents, the information-bearer, and the world.

or

[INF*] Information is the objective structural relationship between logical spaces ([LOGSP*]) and the information-bearer.

It is evident from the conceptualisations I have given, therefore, that I take representation and information to be practically identical. The only noticeable difference between the two conceptualisations, in fact, is the difference between representation-device and information-bearer. But we must ask: What does the difference between representation-device and information-bearer amount to?

A representation-device is the source or vehicle of representation in Suarez (2004) terms. What's more, I have stressed in chapter 2 that a scientific representation-device is a processed artefact that has states or contains within itself entities that present the possible states, characteristics, and configurations of the properties or relations of a target system located in logical space by means of measurement [MEAS]. In chapter 2, I made this claim in the context of a discussion about the content of one particular kind of representation-device: (scientific) models. But now I want to extend this claim to all representation-devices in accord with a broader reading of my definition of measurement ([MEAS]), and the relation between measurement ([MEAS]) and logical space ([LOGSP*]). On this extension, then, a portrait or a teddy bear, for example, will also present the possible states, characteristics, and configurations of the properties or relations of a target system located in logical space by means of measurement [MEAS]. It is just that the target systems in question are the people being painted or a bear, and the logical space in question permits of possible states including standing victoriously or being anthropomorphised and friendly.

Moreover, according to [REP*], representation-devices are integral to the enacting of representation, because they are one of the components that the structural relationship that *is* representation obtains between. And so representation-devices, according to [REP*], will be any object or artefact that is necessary to the enacting of the structural relationship that *is* representation.

Up to this point, I have defined information-bearers predominately by means of examples in order not to be drawn into a debate about the distinction between information and information-bearers. I did say at the beginning of the previous chapter, however, that many theorists have made a distinction between information and information-bearer. For instance, it has been argued that a newspaper is an information-bearer and the news is the information, and this example can be extrapolated to an array of other objects. (Think of the radio waves sent back to earth by the Mars rover, for example, which are surely bearers of the information about such things as the Martian landscape and the chemical composition of the Martian soil, but may not themselves be taken to be information). I also said at the end of the preceding chapter that given an [INF*]-based reinterpretation of any Information-A conception of information, logic should be viewed as both an information-bearer and the thing that makes possible the quantification of information.[83]

From the perspective of an [INF*], then, we can see that information-bearers are those objects and artefacts that have states or contain within themselves entities that make perspicuous features of logical space ([LOGSP*]) by interacting with – i.e. codifying, quantifying etc. – logical space. Logic, for example, is the abstract artefact that reveals features of logical space: this is the very purpose of logical investigations. And, similarly, information-bearers like newspaper articles also display the nature of a logical space by specifying the possibilities that we take to be permissible – i.e. there may be water on Mars. But now, given [INF*] itself, we must also remember that information-bearers contribute to the enacting of information, because information-bearers are one of the components that the structural relationship that *is* information obtains between. So information-bearers, according to [INF*], can be defined as any object or artefact that has states or contain within itself entities that present the structural relationship that *is* information in the process of

[83] As well as a more complete understanding of the additivity, storability, and transmissibility of information.

codifying, describing, or otherwise presenting the nature of logical space ([LOGSP*]).

Putting side by side the definition of representation-device and information-bearer, then, we get the following:

[Representation-device] is any object or artefact that has states or contains within itself entities that present the structural representation that *is* representation in the process of presenting the possible states, characteristics, and configurations of the properties or relations of a target system located in logical space by means of measurement [MEAS].

[Information-bearer] is any object or artefact that has states or contain within itself entities that present the structural relationship that *is* information in the process of presenting the nature of logical space ([LOGSP*]).

As is evident, therefore, the only noticeable difference between representation ([REP]) and information ([INF]) turns out to be little more than a superficiality, because both representation-devices and information-bearers do the exact same thing: enact representation ([REP])/information ([INF]) by presenting logical space in some way (perhaps as a model, a set of logical axioms, or a newspaper article). It follows, then, we find that there is no clear-cut difference between [REP] and [INF].

I argue, therefore, that the concepts of representation and information given by [REP] and [INF] should fall together; that is, that our conceptualisation of representation and information should be merged. What I mean by this is that representation and information should no longer by viewed as two separate and distinct concepts, but rather as two names given to one concept. Thus, I propose that my formulation and comparison of [REP] and [INF] compels a *conceptual synthesis* of the previously separate and distinct concepts representation and information. The conceptual synthesis I am speaking of here is the product of the project of conceptual analysis that I have undertaken thus far in this essay. At its basis, the conceptual synthesis of representation and information I recommend follows from the realisation that the concepts of representation and information are different names given to the way we describe, explain, and communicate the same thing; namely, *an enacted structural relationship between agent and world*. In this way, both the concept of representation and the concept of information can be said to have the same extension: the enacted structural relationship between agent and world [LOGSP*], and some object or artefact. In much that same way as astronomers once recommended that we synthesise our concept of the sun and our concept of the stars, then, I recommend that we synthesise our concepts of representation and information based upon both [REP] and [INF]. For want of a better term, I call the result of this conceptual synthesis: informational representation. And I conceptualise informational representation as follows:

[INFREP] Informational representation is an objective structural relationship between (the semantic activities of) intentional agents; an object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]); and the world.

or

[INFREP*] Informational representation is an objective structural relationship between an object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) and logical space ([LOGSP*]).

Enacting, Presenting, and Making Sense

Perhaps the most obscure aspect of [INFREP*] is that it says that the existence of an object or artefact that presents the possible states, characteristics, and configurations of a logical space, is a necessary condition for the enacting of a structural relationship that *is* informational representation.

It should first be recognised that the object or artefact that presents the possible states, characteristics, and configurations of a logical space – the representation-device or an information-bearer – is often *prima facie* taken to be the representation or information itself. For example, a painting, a collection of ballet dancers, a sculpture, a mathematical model, a newspaper, radio waves, symbols of one form or another etc., may all be taken *prima facie* to be the representation or information itself. But we have seen that this way of thinking cannot be sustained following from a closer analysis of philosophical attempts to conceptualise representation and information. We found, in fact, that the objects or artefacts that present logical space are only one part of the structural relationship we mean to pick out with our concepts of representation and information.

But now the questions arises: What does it mean to say that an object or artefact presents the possible states, characteristics, and configurations of a logical space? And why are these objects and artefacts only a necessary, but not sufficient, condition for the structural relationship that *is* informational representation? To answer these question we must consider what occurs in the process of presenting logical space ([LOGSP*]) and how this presentation of logical space makes possible the enacting of informational representations ([INFREP*]).

In considering the philosophical movement of enactivism at the end of my first chapter above, I mentioned that enacting is usually conceived of in terms of 'the generation of meaning' that occurs when cognitive agents engage 'in transformational […] interactions' (Di Paolo, Rhohde, De Jaegher, 2014: 33).[84] To capture the idea of transformation generation in the context of informational representations, I will use the term *objectification*, whereby the possible states, characteristics, and configurations of a logical space jointly constituted and instantiated by agent and world become presented as objects or artefacts in their own right.[85] Thus, the claim is that where *objectification* occurs and the possible states, characteristics, and configurations of a logical space are presented by objects or artefacts in their own right, then such objects and artefact are able to contribute to the enacting of informational representations

[84] Enactivism is ordinarily taken to be a perspective from which 'epistemology and theories of mind and theories of evolution are very close to being the same thing' (Bateman, 1997: 38). My use of the term enactivism here – and hence my enactivist perspective – is, of course, less about a theory of mind and more about a theory of representation and information. See footnote 29 above for a reminder of how I situate myself in regards to the standard enactivist tradition.

[85] It is important to recognise that I am *not* saying that logical space isn't an object in its own right. Instead, by employing the term objectification I mean to capture what occurs when properties and relation located in logical space become *something more than* mere properties and relation located in logical space; when they are made into or projected into objects and artefacts.

([INFREP]) when they enter into a structural relationship with logical space itself ([LOGSP*]).

I have, in fact, already engaged with the question, what does it mean to say that an object or artefact presents the possible states, characteristics, and configurations of a logical space?, in my discussion of models and measurement in chapter 2. There I argued that models present the properties and relations of a target system located in logical space. In the light of what was said only a moment ago, therefore, it becomes clear that models are *objectifications* – in my sense – of the properties and relations of a target system located in logical space. This objectification can be concrete, as with the foam model of the human cell at NEMO; or it can be abstract, as with a data or mathematical model.[86] All that matters is that the properties and relations that are located in logical space take on a form that makes them very obviously a *presentation of* logical space. This may be as a result of the purpose of which they were created (scientific models are created to explain, describe, predict etc.) or because of a new way of looking at the object or artefact in question (such as making the link between the concentric rings of a tree and the trees age, or coming to appreciate how the carving in a piece of rock describe Egyptian priests, the Macedonian Pharaoh Ptolemy V Epiphanes, and the siege of Lycopolis).

Now, it is quite evident that the presenting of possible states, characteristics, and configurations of a logical space ([LOGSP*]) by

[86] I do not want to commit myself here to the claim that fictional models are also objectifications of possible states, characteristics, and configurations of a logical space ([LOGSP*]). But, in principle, I have no issue with those who would want to defend such a position.

objects and artefacts that function as representation-devices and information-bearers cannot get going without both (the semantic activities of) intentional agents and the world. This is obvious if we realise that the logical space being presented only exists as a consequence of being jointly constituted and jointly instantiated by agent and world ([LOGSP*]). But it is also clear on a more trivial level: you require agents and the world to have such things as paintings, collections of ballet dancers, sculptures, (mathematical) models, newspapers, symbols of one form or another, etc. The presenting of logical space, therefore, is *poetic* – or in my terms the product of explicit *objectification*. And, what's more, the presenting of logical space is *autopoetic*. Why? Because the *objectification* that occurs when possible states, characteristics, and configurations of a logical space ([LOGSP*]) are presented is caused by factors that themselves jointly constitute and instantiate logical space. Put differently, the presenting of logical space through the *objectification* involved in painting, sculpting, modelling, publishing, emitting, writing etc. is a process that is made possible *and* made actual by agent and world.

We can now begin to appreciate why an object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) is a necessary, but not sufficient, condition for the enacting of the structural relationship that *is* informational representation ([INFREP]). Informational representations obtain as a structural relationship between objects or artefacts that present logical space and logical space itself. In order to be differentiable from logical space ([LOGSP*]), however, informational representations must be made actual as *something more* than logical space. That is, informational representations, in order to be differentiable from logical space ([LOGSP*]), must follow from an *objectification* of logical space in my terms. This is why an object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) is a necessary, but not sufficient, condition for informational representations: because without such an object it would be impossible to differentiate logical space (a collection of possibilities) from informational representations (an enacted actuality): we would have only logical space whereby we could only specify what representation and information are about in terms of possibilities.

But then one could ask: Why not just say that the object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) *is* the informational representation?

Objectification, as I have already said, can take the form of a painting, a collection of ballet dancers, a (mathematical) model, a collection of symbols, and a whole host of other things besides. But if we are to say that these objects or artefacts *are* informational representations in virtue of presenting the possible states, characteristics, and configurations of a logical space ([LOGSP*]), then we encounter some problems. Think again of the Rosetta Stone/Egyptian hieroglyphics example. Certainly for the Ancient Egyptians who created the Rosetta Stone the hieroglyphics carved into pyramid walls or marked onto pieces of papyrus were instances of informational representations. But if we say that these objects that present logical space ([LOGSP*]) – i.e. carved stones, marked papyrus etc. – *are* informational representations, then what are we to say about these objects in the eighteenth century prior to the point in time when the Rosetta Stone had been deciphered? Did they count as informational representations then?

I have already said that on the [INF*] conceptualisation of information the Rosetta Stone would not count as information independently, because the Rosetta Stone (or a sheet of papyrus, for that matter) as an information-bearer is a necessary, but not sufficient, condition for information. And my reasoning for that claim can now be used to explain why it will not do to say that any object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) *is* the informational representation. The reason is simple: because an informational representation must always obtain *between* agent and the world such that the agent is able to make sense of the way that an object or artefact presents or *objectifies* the possible states, characteristics, and configurations of a given logical space ([LOGSP*]). That is, informational representations are always enacted equally by representation-device/information-bearer, agent, and world. Thus, without someone around to make sense of the way that an object or artefact presents or *objectifies* the possible states, characteristics, and configurations of a given logical space ([LOGSP*]), we can have no instance of informational representation ([INFREP*]). But why, then, is this *making sense* of such importance? And why does the need for making sense explain why it will not do to say that any object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) *is* the informational representation.

To make sense of something an agent or group of agents must be able to identify what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) an object presents or *objectifies*. For example, in the mathematical models given by Einstein's theory of relativity an agent must be able to make sense of the way that certain symbols – i.e. c – present or *objectify* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) – i.e. a logical space that allows for the possibility that the speed of light in a vacuum is 299 792 458 m/s and can be no greater. And in the case of Pablo Picasso's painting, *Guernica*, an agent must be able to make sense of the way that certain marks on the canvas present or *objectify* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) – i.e. a logical space that allows for the possibility of people, animals, and buildings suffering as the result of violence and chaos of some sort (perhaps by the bombing of a Spanish country village, for instance). And the important point is as follows: *without making sense of how an object presents or objectifies possible states, characteristics, and configurations of a given logical space ([LOGSP*]) that object cannot be said to be a representation-device/information-bearer*. Why? Because the [INFREP] conceptualisation is designed to be non-dualistic; that is, to privilege neither agent nor world. However, if we allow that an object can present or objectify possible states, characteristics, and configurations of a given logical space ([LOGSP*]) without anyone to present or *objectify* too, then we fall back into a conception of information and representation that discounts the semantics and intentionality of agents. In this way, we would implicitly advocate an environmental conception of information and, potentially, a relational conception of representation; and this, I think, would be to open the door once again to philosophical dualism.

Bringing this discussion back to the Rosetta Stone/Egyptian hieroglyphics example, it becomes clear that the hieroglyphics would only count as informational representations ([INFREP]) if an agent were able to make sense of the way that the hieroglyphics present or *objectify* possible states, characteristics, and configurations of the logical space of the Ancient Egyptians ([LOGSP*]). And this is exactly what eighteenth century man could not do, because without the Rosetta Stone it was unclear how the hieroglyphics should be read (if at all), and hence was unclear how (if at all) the hieroglyphics presented or *objectified* possible states, characteristics, and configurations of the Ancient Egyptian's logical space ([LOGSP*]). Just how eighteenth century man learnt to make sense of the hieroglyphics is a complex story involving such things as military engagements inspired by socio-economic based conflicts (between England and France, in this particular example), Ancient Greek and Demotic texts, and eighteenth century archaeologists and linguists. But what is clear is that making sense of an object that presents or *objectifies* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) by an agent or group of agents is required first before an informational representations itself ([INFREP]) can be enacted.

We have seen, therefore, that whilst an object that presents or *objectifies* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) is a necessary condition for informational representations ([INFREP]), such an object is not equivalent to informational representations. The reason for this is because such an object needs to be made sense of before it can be taken as a representation-device or information-bearer, and hence before it can contribute to the enacting of an informational representation. It is only in the making sense of the object, then, that such an object is able to enter into the relationship with agent and world (i.e. logical space ([LOGSP*]) that enacts an informational representation ([INFREP]).

Presenting, Making Sense, and Accuracy

The notion of making sense of an object that presents or *objectifies* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) is still vague for us at this point. We may ask, for instance: What is making sense and how does such making sense occur? I propose that the notion of making sense can be grasped most easily by means of examples.

Consider the exhibits and exhibit descriptions at NEMO science museum in Amsterdam that I introduced in the introduction of this essay. In the case of those objects, it is clear that individuals who are unfamiliar with the relevant theories in genetics, biology, engineering, or physics will have to make sense of the way that the exhibits and exhibit descriptions present or *objectify* the possible states, characteristics, and configurations of a given logical space ([LOGSP*]) if they want the exhibits and exhibit descriptions to contribute to the enacting of informative representations ([INFREP]); i.e. if they want the exhibits and exhibit descriptions to be informative representations *for them*. This making sense will ordinarily involve the interpretation of an object through the lens of a pre-mastered syntactical and semantic framework, where both syntax and semantics are not to be understood in exclusively linguistic terms, but rather as the general rules and meaning by which a chosen system, group, code, or language is formed, composed, and organised. It is clear that to make sense of an exhibit descriptions at NEMO one merely has to read what is has been written on the label accompanying the exhibit (so long as one is linguistically competent). But for the exhibits themselves making sense will involve either (a) the interpretation of the exhibit through the lens of a pre-mastered syntactical and semantic framework that one takes to be helpful in identifying what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) the exhibit in question presents or *objectifies*; or (b) the interpretation of the exhibit through the lens of a pre-mastered syntactic and semantic framework that was suggested by one's making sense of another representation-device/information-bearer (i.e. the exhibit description) to identify what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) the exhibit in question presents or *objectifies*. The problem with (a) is that one is not guaranteed to make sense of the exhibit as one 'should' – as in the case of the child who takes the foam model exhibit of the human cell to be a foam model of the eye of some exotic animal. And the problem with (b) is that there are not always accompanying representation-devices/information-bearers that one can easily make sense of; i.e. there are not always other objects to guide how we are to make sense of any given object or artefact.[87] [88]

[87] I think the problem with (b) is, in fact, the problem contemporary scientists and philosophers have with interpreting quantum mechanics (QM). The reason being that we are struggling to make sense of the objects and artefacts (mathematical models) associated with QM primarily because there are no other available objects and artefacts that can be easily made sense of to help us specify unequivocally what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) are presented or *objectified* by the objects and artefacts (mathematical models) associated with QM themselves. This state of affairs in reflected in the multitude of differing interpretations of QM and, most obviously, in Niels Bohr's Complementarity Principle, which states that the classical concepts are necessary to give a physical description of the world even if QM descriptions are incommensurable with the classical descriptions (Bohr, 1935 and 1959; Ozawa, 2003).

[88] It should be said that a person who creates an object for him/herself will automatically know what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) that object presents. But, still, the possibility of being able to create a representation-device/information-bearer requires the mastery of a syntactical and semantic framework in the first place. This, I think, is one of the most

Making sense, therefore, is essentially a process of projecting a set of semanticised rules onto an object or artefact to make that object or artefact – or, more precisely, the possible states, characteristics, and configurations of a given logical space ([LOGSP*]) presented or *objectified* by that object or artefact – tractable. And so (a) and (b) above can be taken to be different methods for how one could go about projecting a set of rules onto an object or artefact. An interesting consequence of this observation, I think, is that art can then be understood as the domain in which the rules that can be projected onto an objects and artefacts are most open to interpretation. This is less the case when the artwork is a realist portrait or sculpture, because then the rules one is supposed to project onto the artwork to make it tractable are more fixed. But now think of an expressionist painting, an abstract piece of dance, or a piece of modern art involving a urinal. It may be not be clear which rules one is "supposed" to project in these examples at all, and so any number of different rules may be projected onto these objects and/or artefacts to make tractable the possible states, characteristics, and configurations of a given logical space ([LOGSP*]) that the objects *may be* presenting or *objectifying*.

Crucially, however, if the object or artefact is to contribute to the enacting of an informational representation ([INFREP]) at all (and hence if the object or artefact is to be a piece of art at all, on my view), then it must present or objectify some possible states, characteristics, and

important skills children learn in their development. Put differently, I think that the capacity to create and *make sense* of objects and artefacts that present or objectify possible states, characteristics, and configurations of a given logical space ([LOGSP*]) – and hence the possibility to enact informational representations ([INFREP]) – is one of the key features of a child's education.

configurations of a given logical space ([LOGSP*]), even if it is difficult to make sense of what possible states, characteristics, and configurations of a given logical space the object or artefact is presenting or *objectifying*. This all implies that the making sense of an object or artefact must be possible if the object or artefact is to count as having presented or *objectified* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) at all. That is, if the object or artefact is to count as one of the components that enters into the structural relationship that *is* informational representation [INFREP].

But then the question arises: How important is accuracy to our making sense of an object or artefact? Certainly, we would feel obliged to correct the child in NEMO who makes sense of the foam model of the human cell in terms of an eye of some exotic animal. But would we feel obliged to correct someone who made sense of a piece of art in a way that was different to how we ourselves made sense of that piece of art? Or would we correct a child who made sense of a four poster bed in terms of an aeroplane or a stuffed piece of material in terms of a friendly bear? The question I am probing at the moment is the following: How can we know when one particular way of making sense of an object or artefact is to be favoured over another?

One way to assess which ways of making sense of an object or artefact are to be favoured is to ask whomever it was that made the object or artefact what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) the object or artefact is supposed to present or *objectify*. This method, however, will only work when the creator of the object or artefact is available to ask and knows for him/herself what possible states, characteristics, and configurations of a given logical space ([LOGSP*]) the object or artefact is supposed to present or *objectify*. We may think that the creator of an object or artefact will always be clear on what the object presents or *objectifies*, but nothing could be farther from the truth. Take modern mathematical physics, for example, where mathematical models are devised according to principles of, for instance, simplicity (Occam's Razor), consistency, completeness, elegance, explanatory power, etc. The physicists who devise such objects/artefacts mathematical models according to these principles are often quite unclear about how to interpret such objects/artefacts and so are not themselves authorities on how to make sense of such objects/artefacts. What's more, some objects or artefacts that seem to present or *objectify* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) do not appear to have been created by anyone at all. Consider the series of concentric rings visible in the wood of a cut tree trunk or the cosmic background radiation. One can make sense of these objects as presenting or *objectifying* possible states, characteristics, and configurations of logical spaces ([LOGSP*]) concerned with tree ages, big bangs, left over heat, initial conditions of the universe etc., but any question about what is the right way to make sense of such objects could only ever be directed at nature itself.[89]

[89] It is now widely accepted that human beings have a tendency to assume that we have found in reality a great number of objects that are presenting or *objectifying* possible states, characteristics, and configurations of a given logical space ([LOGSP*]). Consider the familiar experience of seeing face in the landscape of the moon or in the headlights and grill of an automobile. Such pattern recognition – or *apophenia* – may tend to make us believe in the existence of objects that present or *objectify* possible states, characteristics, and configurations of other logical spaces ([LOGSP*]) (an example would be divination in which an object – tea leaves, for instance – is taken to present or *objectify* possible states, characteristics, and configurations of a logical spaces

The conclusion is that in many cases we are unsure how best to make sense of objects and artefacts that we can – both in theory and in practice – take to be presenting or *objectifying* possible states, characteristics, and configurations of a given logical space ([LOGSP*]). In certain domains this is unproblematic. In art and child's play, for example, the value of making sense of which possible states, characteristics, and configurations of a logical space ([LOGSP*]) an object or artefact presents or *objectifies* is usually determined instrumentally by asking which way of making sense let's one best appreciate a piece of art, or which way of making sense allows one to have fun whilst playing (the snoopy cartoon in the front section of this essay is an allusion to this last remark). But in the domain of science the uncertainty as to how best to make sense of objects and artefacts that purportedly present or *objectify* possible states, characteristics, and configurations of a given logical space ([LOGSP*]) is problematic. Why? Because if we cannot decide how best to make sense of an object or artefact in this regard, then we have no criteria for determining which way of making sense of an object or artefact is the most accurate. And this opens the door to a pernicious form of relativism that one may feel inclined to reject or, at the very least, avoid.

Recall, then, that [INFREP] holds that informational representations are enacted by a tripartite relational structure involving (the semantic activities of) intentional agents; an object or artefact that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]); and the world. As we have seen, we may have some difficulty with deciding how we are to make sense of the object or artefact

concerned with future events in someone's life). The concern again in this instance is how we are to decide how – if at all – we should be making sense of such objects.

that presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]). This leads to a problem that can be stated as follows. If we are unsure how best to make sense of an object that purportedly presents or *objectifies* logical space, how then could we know if the informational representation ([INFREP]) that *is* the enacted structural relationship between such an object, agent, and world is ever accurate? And, relatedly, how then can we know where and when any given enacted informational representation ([INFREP]) should be maintained and endorsed?

Informational Scientific Representations

At this stage, we must recall my conceptualisation of scientific representation [SCIREP] given in chapter 2. I defined scientific representation as follows:

[SCIREP] A scientific representation is the structural relationship between scientifically admissible logical spaces ([LOGSP*]) and the scientific model.

At that point, I argued that scientific representations can be differentiated from other kinds of representation, because the logical spaces that are presented by scientific models (i.e. scientific representation-devices) are constrained in the sense that they can only permit scientifically admissible possible states of affairs. And thus the enacted structural relationship between scientifically admissible logical space and scientific models is always subject to the constraints of scientific admissibility (see pages 90-93 above for a reminder of the argument).

I also said that scientific admissibility is fixed by institutional norms and error filters that – according to my [SCI] conceptualisation – are *aimed* at discovering the objective character of the world. Thus, the constraints of scientific admissibility on the enacting of scientific representation can be interpreted as constraints that regulate the production of scientific representation-devices in one important respect: they make it the case that scientific representation-devices must presents the possible states, characteristics, and configurations of a logical space ([LOGSP*]) that is at least aimed at giving the nearest approximation of reality possible.

Now, returning to the discussion of our making sense of objects and artefacts that present or *objectify* logical space, we can engage in one final discussion before reaching the crescendo of this chapter. Any object can be made sense of in some respect. Cloud patterns in the sky provide the perfect example here, but literally any object, collection of objects, artefact, or collection of artefacts will suffice. The number of possible ways one can make objects and artefacts tractable will always be very large. In fact, the number of possible ways of making sense of an object and artefact will be as large as the syntactic and semantic framework one has mastered. So if one has mastered a greater number of syntactic and semantic frameworks, then one will be able to project onto an object or artefact a greater number of semanticised rules for making that object or artefact tractable. In the case of clouds, one will be able to make sense of the object in terms of a map of British Honduras on the Caribbean, the profile of Thomas Eakins, the stoning of Stephen, or a ducky or as horsie.[90] And by making sense of this object in whatever way – and hence

[90] I am referring once again to the Charlie Brown comic strip I presented in the front section.

taking the object to present or objectify certain possible states, characteristics, and configurations of a logical space ([LOGSP*]) and not others – one directly influences the kind of informational representation ([INFREP]) that is to be enacted between the object, agent, and world.

We have seen in the case of scientific representations, however, that scientific representation-devices (models) can only present or objectify possible states, characteristics, and configurations of scientifically admissible logical space ([LOGSP*]) if they are to count of scientific at all. So the objects and/or artefacts that are scientific representation-devices *should* be made sense of by consideration of the constraints of scientific admissibility. That is not to say that one cannot make sense of a scientific representation-devices differently. A child visiting NEMO is perfectly capable of making sense of the foam model of the human cell in terms of the eye of a monster by projecting the semanticised rules she has mastered concerning monsters and eyes onto the model. But in the case of scientific representation-devices one *should not* make sense of the object and/or artefact in this way, because to do so would contravene the dictates of scientific admissibility that constrained the way that the object and/or artefact in question is supposed to present logical space. And so it would undermine the [SCI]-based claim that if we follow the dictates of scientific admissibility we are able to discover to an ever greater approximation the nature of the world.

It should be noted at this stage that the constraints of scientific admissibility are fixed by the latest conclusions supported by the institutional norms and error filters of science [SCI], and not by one's own knowledge of what is permitted by these institutional norms and error filters. If it were fixed by one's knowledge of what is permitted by these institutional norms and error filters in question, then it would follow that a child who knew nothing about human cells but everything about ophthalmology may be justified in making sense of the foam model in NEMO in accordance with her best knowledge of a scientifically admissible logical space ([LOGSP*]) centred on eyes and not human cells. Of course, we would all agree that making sense of the object in this manner is to be considered as more scientific than making sense of the object as something from a fantasy fiction novel. But accuracy must be indexed to the institutional norms and error filters of science [SCI] independent of one's own awareness of the latest conclusions supported by such institutional norms and error filters. That is, the accuracy of a given object must be determined relative to the intersubjective institutional norms and error filters of science [SCI], and not one's own understanding of these institutional norms and error filters. The central purpose of our education, I think, is to make us aware to an ever greater extent about the possible states, characteristics, and configurations of a logical space ([LOGSP*]) that is appropriately constrained by scientific institutional norms and error filters. Our education, then, is directed towards a more complete understanding of scientific admissibility, where scientific admissibility is equivalent to accuracy.

Given what I have been saying, therefore, it is evident that if we want to make sense of an object that purportedly presents or *objectifies* logical space in the most accurate way possible, then such objects *should* be made sense of by consideration of the constraints of scientific admissibility. Thus, we have a two-way corrective mechanism of accuracy. Firstly, the object itself should present or objectify only scientifically admissible logical space – this is a constraint on the construction of the object in question. And, secondly, one should make sense of the object by consideration of the constraints of scientific admissibility that the object should adhere to when presenting or objectifying logical space – this is a constraint on the use of the object in question. If either of these two checks fail, then the accuracy of the object is negatively affected. Of course the standard of accuracy will vary in accordance with variations in scientific admissibility. But it is still the case that the most accurate scientific representation-devices at time *t* will be those that follow the dictates of scientific admissibility at time *t*. The accuracy of a scientific representation-device, then, can be assessed by evaluating how far the representation-device is consistent with the dictates of scientific admissibility at any given point in time.

At this point, we can finally transpose what I have been saying about scientific-admissibility and scientific representation-devices onto the discussion about informational representations ([INFREP]). We can do this by recalling two claims I have made in this chapter. Firstly, that both representation-devices and information-bearers are equivalent in the sense that what they both do is present logical space. And, secondly, that it is only by making sense of how an object presents logical space that such an object is able to enter into the relationship with agent and world – i.e. logical space – that enacts an informational representation ([INFREP]). It follows, therefore, that the evaluation of how far a representation-device is consistent with the dictates of scientific admissibility will also be an evaluation of how far the informational representation ([INFREP]) is consistent with the dictates of scientific admissibility. This is the case because the corrective mechanisms of accuracy I have just described apply equally to all those objects that contributed to the enacting of informational representations ([INFREP]) by presenting logical space. Thus, informational representations that are enacted in part by objects or artefacts (representation-devices/information-bearers) that are consistent with the dictates of scientific admissibility will be more accurate. In fact, I propose that such informational representations will be of a qualitatively different kind than other informational representations, because they will be guided by an attempt to give the best approximation of the world possible. So scientific admissibility will determine which informational representations should and should not be maintained and endorsed relative to the aim to give the best approximation of the world.[91] I call informational representations that are bound by this aim to give the best approximation of the world: *informational scientific representations.*

[INFSCIREP] Informational scientific representation is an objective structural relationship between (the semantic activities of) intentional agents (including *making sense* in accordance with scientific admissibility); an object or artefact that presents the possible states, characteristics, and configurations of a scientifically admissible logical space ([LOGSP*]); and the world.

Or

[INFSCIREP*] Information scientific representation is an objective structural relationship between an object or artefact that presents the possible states, characteristics, and configurations of a logical space in

[91] This claim does not hold for informational representations in the domain of art or child's play, however.

accordance with the dictates of scientific admissibility (i.e. a scientific model) and a scientifically admissible logical space ([LOGSP*]).

Concluding Remarks: Perspectivism in Philosophy of Science and Information

In this essay, I first reconceptualised representation as a structural relationship. Then I differentiated scientific representation from other kinds of representation. I moved on to give a reconceptualisation of information as a structural relationship. Before arguing that we should synthesise our concepts of representation and information into one concept: informational representation. Finally, I argued that we can differentiate informational scientific representation from other kinds of information representation. Now, I will draw some final conclusions in regards to my introduction of informational (scientific) representations and provide answers to some of problems that remain open.

We can approach the most important of these problems by asking: What is it, exactly, that informational representations ([INFREP]) are representing and informing us about? The answer to this question is as follows: informational representations ([INFREP]) are representing and informing us about the multi-layered relationship between the world, objects and artefacts in the world, and agents who use objects and artefacts to think about the world. But, then, one may ask: Are we or are we not representing or being informed about the world in-itself? The answer to this question is both yes and no, because we are representing and being informed about the world *as it stands in a relationship to agents who use objects* *and artefacts to think about the world*. The world that we take ourselves to be representing or being informed about, therefore, is mediated through the relationship between agent, object and artefact, and world (as manifest in logical space and the presenting of logical space ([LOGSP*])). It is only a remnant of philosophical dualism that makes us think that the world can be conveniently isolated and studied as an independent – i.e. agent and object/artefact free – entity.

At the end of the first chapter, I noted that conceptualising representation as a structural relationship could cause us some problems when we come to think about past representations that have turned out to be representations of nothing; i.e. phony representations. The example I used then was the models, theories, and laws of phlogiston from the eighteenth/nineteenth century. My conceptualisation of informational (scientific) representation, however, allows for us to approach the question, are the models, theories, and laws of phlogiston to be considered genuine representations or not?, from a new perspective (see pages 52-53 above for a reminder of the problem). The conceptualisation of informational (scientific) representation ([INFSCIREP]) I offer makes it possible to say that the models, theories, and laws of phlogiston were – and are! – informational representations ([INFREP]) even if we now hold that phlogiston does not exist. Importantly, however, they are no longer informational *scientific* representations ([INFSCIREP]). The reason for this is because the models, theories, and laws of phlogiston taken as informational representation *simply represent or inform us about a world as it stands (or stood) in a relationship with agents who make (or made) use of certain objects and artefacts to think about the world*. This is the case because the relationship between a particular group of agents (scientists/physicists) and world during the eighteenth and nineteenth centuries was such that the models, theories, and laws of phlogiston did represent and inform us about something (i.e. phlogiston) that was at least taken to exist. In other words, the models, theories, and laws of phlogiston were and are genuine informational representations ([INFREP]), because the representation of, and information about, phlogiston was brought about as the result of the relationship between eighteenth and nineteenth century agents, the objects and artefacts they used, and the world.[92] Moreover, the models, theories, and laws of phlogiston can also be said to have been genuine informational *scientific* representations ([INFSCIREP]) so long as the purported phlogiston being represented and being informed about did not contravene the dictates of scientific admissibility *when enacted in the eighteenth and nineteenth centuries*.

Now, it is evident that today the models, theories, and laws of phlogiston could not be said to be informational *scientific* representations ([INFSCIREP]), because they contravene the dictates of contemporary scientific admissibility. But that is not to say that they are not still informational representations ([INFREP]) in the sense that what they have the potential to represent and inform us about is a relationship between agent and world that we could adopt (I can think that phlogiston

[92] This brings us neatly back to the problem we faced at the end of chapter 1 about [REP]'s conceptualisation of the nature of the target of a representation-device (source). There we were concerned with the question, does the reality or non-reality of a target system influence whether or not the structural relationship that *is* representation obtains? At this point, however, we can see that in the case of representation (and information) the target system in question is always given by means of a mapping in logical space, which is then *objectified* by the source (i.e. the representation-device/information-bearer). This source is then to be made sense of, where the accuracy of the source (and hence the mapped target system) *can only be assessed relative to the standards and dictates of scientific admissibility at any given time t*.

exists if I really want to). Thus, just in the same way as we can think that our arms are wings or that a bottle of shampoo in the shower is an Oscar award, we can also think that phlogiston is that which causes combustion. What's more, we can present the possible states, characteristics, and configurations of a phlogiston-permitting logical space ([LOGSP*]) by *objectifying* that logical space in a model, theory, or law of phlogiston. And, then, an informational representation ([INFREP]) will be enacted between this object/artefact and logical space that represents and informs us about a relationship between agent and world in which phlogiston is taken to exist.

Informational representations, therefore, emerge from our capacity to think about the multiplicity of ways the world *could* be. That is, different informational representations are made possible because by altering our relationship with the world we are able to consider different possible states, characteristics, and configurations of the world. This just is what it means to work in different logical spaces, because logical space is jointly constituted and jointly instantiated by agent and world: logical space – and with it the multiplicity of ways logical space can be – *is* the relationship between agent and world. Our capacity to work in different logical spaces – i.e. to take up different standpoints to the world – is, then, a consequence of our manipulating the relationship we have with the world.[93] And informational representations should be taken as the presentation – or *objectification* – of these standpoints. This is why the presenting of logical space by an object is so important, because it displays or exhibits the logical space being presented and reveals what possible states,

[93] Any time you have read a piece of fiction or watched a movie you have engaged in this kind of manipulation.

characteristics, and configurations are permitted by the logical space in question. This also explains why it is imperative that one is able to make sense of what logical space an object actually presents; i.e. represents and informs about (think back to the Rosetta Stone example here).

But now, of course, one may want to query: how can we measure the truth of informational representations? The answer is that the truth of informational representations is not to be found in a correspondence with the world or coherence of our beliefs, but rather in the shaping the relationship between agent and world by means of mediating objects and artefacts. Truth, then, is relative to the relationship between agent, objects or artefacts, and world, and can only be fixed to the nature of this relationship. This statement is directly analogous to my remarks about the accuracy of informational representations in the last chapter. The truth of an informational representation, then, like accuracy, is not something fixed by the world (correspondence) nor is it something fixed by the semantic activities of the agent (coherence). Rather, truth – if it is to be found anywhere – is to found in the dynamic relationship through which agents use objects and artefacts to think about the world. This relationship, as I have said over and over again by this point, *is* informational representation, which in turn *is* the relationship between logical space and the *objectifying* objects and artefacts that present logical space.

But where, then, does this leave us in regards to informational scientific representations ([INFSCIREP])? I have said already that informational scientific representations are bound by the aim to give the best approximation of the world possible. But is this anything more than an aim? Do informational scientific representations really give the best –

i.e. the most truthful – approximation of the world? This is *the* question that underpins the debate between scientific realists and scientific anti-realists.

My answer is that we should think about informational scientific representation – and, as a consequence, information, representation, and science – in perspectival terms. In this way, I follow Giere's (2006) claims that scientific models represent the world from a particular vantage point or point of view. Giere says:

> all theoretical claims remain perspectival in that they apply only to aspects of the world and then, in part because they apply only to some aspects of the world, never with complete precision.

And, moreover, that:

> a particular scientific perspective arises out of both social interactions among members of a scientific community and interactions with the world, typically mediated by complex instrumentation. (Giere, 2006: 15).

Thus, I hold with Giere that truth claims are 'relative to a perspective', where for perspective we can read the relationship between agent, object or artefact, and world in my terms (Ibid.: 81). Scientific admissibility can, then, be understood as a constraint on scientifically viable perspectives. Just as *objectifying* objects and artefacts can be understood as that which presents or makes manifest a given perspective.

Interestingly, therefore, given what I have just said about perspectivism, this implies that the only way to resolve the debate between scientific realists and scientific anti-realists would be to get clear on the nature of the ever-changing relationship between agent(s), object(s) or artrefact(s), and world that *is* informational scientific representation. And this is just to say that to decide between scientific realism and scientific anti-realism we must get clear on the nature of my [REP] and [INF] conceptualisations of representation and information. To do this we must, I think, focus on two key aspects of both representation, information, and the conceptual synthesis of the two I endorse, informational (scientific) representation: their defeasibility and their shareability. We must get clear, that is, on the nature of representation, information, and thus informational (scientific) representation that makes them (a) open to being annulled or invalidated, and (b) intersubjectively communicable.

This research will have to be left for the future. But it is imperative now that I make clear the logical consequences of my perspectival account of informational scientific representations as a dynamic relation between agent(s), object(s) or artefact(s), and world. The central idea is that on my view to decide between scientific realism and scientific anti-realism; to give an account of the purpose and nature of representation; and to give an account of the purpose and nature of information, all fall together. In this way, the questions, does science give a true description of the world?, what is the nature of representation?, and, what is the nature of information?, cannot be distinguished. Thus, I propose that now disparate domains of philosophy – i.e. the philosophy of science, information, representation – can be brought into contact around one common question that should guide future research: What is the nature of the dynamic relationship between agent(s), object(s) or artefact(s), and world? Or, put differently, what is the nature of informational (scientific) representation?

Bibliography

Adriaans, P.W. & van Benthem, J.F.A.K. (eds.), 2008b, *Handbook of Philosophy of Information*, (Amsterdam: Elsevier Science Publishers).

Adriaans, P., Information, *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.).

Barbrousse, A. & Ludwig, P., 2009, Fictions and Models, In M. Suárez (ed.): *Fictions in Science: Philosophical Essays on Modelling and Idealisation*, (London: Routledge): 56–75.

Bar-Hillel, Y. & Carnap, R., 1953, An Outline of a Theory of Semantic Information, repr. in Bar-Hillel, 1964, *Language and Information: Selected Essays on Their Theory and Application*, (Reading, Mass; London: Addison-Wesley): 221-74.

Bartels, A., 2006, Defending the Structural Concept of Representation, *Theoria*, 55: 7-19.

Barwise, J. & Perry, J., 1983, *Situations and Attitudes*, (Cambridge, MA: MIT Press).

Barwise, J. & Seligman, J., 1997, *Information Flow: The Logic of Distributed Systems*, (Cambridge: Cambridge University Press).

Berkeley, G., 1948–1957, *The Works of George Berkeley, Bishop of Cloyne*. A.A. Luce and T.E. Jessop (eds.). (London: Thomas Nelson and Sons), 9 vols.

Bohr, N., 1935, Quantum Mechanics and Physical Reality, *Nature* 136: 1025-1026.

–––, 1959, Discussion with Einstein on epistemological problems in atomic physics, in *Albert Einstein: Philosopher-Scientist*, P. A. Schillp (eds.), (New York: Harper and Row).

Braithwaite, R. B., 1953, *Scientific Explanation: A Study of the Function of Theory, Probability and Law in Science*, (Cambridge: Cambridge University Press).

Braman, S., 1989, Defining Information, *Telecommunications Policy,* 13: 233-242.

Bueno, O., 1997, Empirical Adequacy: A Partial Structures Approach, *Studies in History and Philosophy of Science*, 28: 585–610.

Cartwright, N., 1989, *Nature's Capacities and their Measurement*. (Oxford: Oxford University Press).

Chaitin, G., 1987, *Algorithmic Information Theory,* (Cambridge: Cambridge University Press).

Chakravartty, A., 2007, *A Metaphysics for Scientific Realism*, (Cambridge: Cambridge University Press).

–––, Informational versus Functional Theories of Scientific Representation, *Synthese* 72: 197-213.

Checkland, P. B., & Scholes, J., 1990, *Soft Systems Methodology in Action* (New York: John Wiley & Sons).

Clark, A., & Toribio, J., 1994, Doing without representing, *Synthese*, 101: 401–434

Contessa, G., 2007, Scientific Representation, Interpretation, and Surrogative Reasoning. *Philosophy of Science*, 74 (1): 48-68.

–––, *Representing Reality: The Ontology of Scientific Models and Their Representational Function*, Dissertation, University of London.

da Costa, N. C. & French, S., 2000, Models, Theories, and Structures: Thirty Years On, *Philosophy of Science* 67, Supplement: S116–127.

–––, 2003, *Science and Partial Truth: A Unitary Understanding of Models and Scientific Reasoning*, (Oxford: Oxford University Press).

Davis, G. B., & Olson, M. H., 1985, *Management Information Systems: Conceptual Foundations, Structure, and Development*, 2nd edn (New York: McGraw-Hill).

Debons, A., & and Cameron, W. J. (eds.), 1975, *Perspectives in Information Science* : *Proceedings of the Nato Advanced Study Institute on Perspectives in Information Science, Held in Aberystwyth, Wales, Uk, August 1324, 1973*, (Leiden: Noordhoff).

deVries, W. A., 2011, Sellars and the Myth of the Given, In M. Bruce & S. Barbone (eds.), *Just the Arguments: 100 of the Most Important Arguments in Western Philosophy*, (London: Wiley-Blackwell).

Devlin, K. J., 1991, *Logic and Information*, (Cambridge: Cambridge University Press).

Di Paolo, E. A., Rhohde, M., De Jaegher, M., 2014, Horizons for the enactive mind: Values, social interaction, and play, in J. Stewart, O. Gapenne, E. A. Di Paolo, (eds.) *Enaction: Toward a New Paradigm for Cognitive Science*. MIT Press: 33.

Dretske, F., 1981, *Knowledge and the Flow of Information*, (Cambridge, MA: The MIT Press).

Dunn, M., and Hardegree, G., 2001, *Algebraic Methods in Philosophical Logic*, (Oxford: Clarendon Press).

Elgin, C., 1996, *Considered Judgement*, (Princeton: Princeton University Press).

–––, 2004, Denying a Dualism: Goodman's Repudiation of the Analytic/Synthetic Distinction, *Midwest Studies in Philosophy*, 28 (1): 226–238.

–––, 2009, Exemplification, Idealization and Understanding, in M. Suárez (Ed.), 2008: 77-90.

Fodor, J. A. 1987, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, (Cam., Mass.; London: MIT Press).

Floridi, L., 2004c, Outline of a Theory of Strongly Semantic Information, *Minds and Machines,* 14 (2): 197-222.

2011, *The Philosophy of Information*, (Oxford: Oxford University Press).

——, 2011a, Semantic Information and the Correctness Theory of Truth, *Erkenntnis*, 74 (2):147-175.

Fraassen, B. C. van, 1980, *The Scientific Image*, (Oxford: Oxford University Press).

——, 1989, *Laws and Symmetry*, (Oxford: Oxford University Press).

——, 2008, *Scientific Representation: Paradoxes of Perspective*, (Oxford: Oxford University Press).

French, S., 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: 1472-1483.

——, 2014, *The Structure of the World: Metaphysics and Representation*, (Oxford: Oxford University Press).

French, S., & Ladyman, J., 1999, Reinflating the Semantic Approach, *International Studies in the Philosophy of Science*, 13: 99–117.

Frigg, R., 2006, Scientific representation and the semantic view of theories. *Theoria. an International Journal for Theory, History and Foundations of Science*, 21 (1): 49-65. ISSN 0495-4548

——, 2010a, Fiction in Science, In: J. Woods (ed.): *Fictions and Models: New Essays*, (Munich: Philosophia Verlag): 247–287.

——, 2010b, Models and Fiction, *Synthese*, 172 (2): 251–268.

Frigg, R., & Hartmann, S. Models in Science, *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), E. N. Zalta (ed.).

Gibbs, J.W., 1906, *The scientific papers of J. Willard Gibbs in Two Volumes*, 1, (Longmans, Green, and Co).

Giere, R., 1991, Explaining Science: A Cognitive Approach. *Philosophical Review*, 100 (4): 653-656.

–––, 1994, No Representation Without Representation, *Biology and Philosophy*, 9 (1): 113-120.

–––, 2004, How Models are Used to Represent Reality, *Philosophy of Science*, 71: 742-752.

–––, 2006, *Scientific Perspectivism*, (Chicago: University of Chicago Press).

–––, 2009, Why Scientific Models Should Not be Regarded as Works of Fiction, In: M. Suárez (ed.): *Fictions in Science. Philosophical Essays on Modelling and Idealisation* (London: Routledge): 248–258.

Godfrey-Smith, P., 2006, The Strategy of Model-based Science, *Biology and Philosophy*, 21: 725–740.

–––, 2009, Models and Fictions in Science, *Philosophical Studies*, 143: 101–116.

Goodman, N., 1976, *Languages of Art: An Approach to a Theory of Symbols*, 2nd edition, (Indianapolis: Hackett Publishing Company).

Graham, G., 1999, *The Internet: A Philosophical Inquiry*, (London: Routledge).

Grice, H. P., 1989, *Studies in the Way of Words*, (Cam., Mass.: Harvard University Press).

Grünwald, P.D., 2007, *The Minimum Description Length Principle*, (Cam., Mass.: MIT Press).

Harremoës, P. & Topsøe, F., 2008, The quantitative theory of information, in Adriaans and van Benthem, 2008b.

Harris, T., 2003, Data Models and the Acquisition and Manipulation of Data, *Philosophy of Science*, 70: 1508–1517.

Herken, R., 1995, *The Universal Turing Machine – A Half-Century Survey*, (Springer Verlag).

Hintikka, J. & Suppes, P. (eds.), 1970, *Information and Inference*, (Dordrecht: D. Reidel).

Hughes, R. I. G., 1996, Models and Representation, *Philosophy of Science*, 64 (4): 336.

Kemeny, J., 1953, A Logical Measure Function, *Journal of Symbolic Logic*, 18: 289–308.

Kratzer, A., 2014, Situations in Natural Language Semantics, *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.)

Kripke, S. A., 1959, A Completeness Theorem in Modal Logic. *Journal of Symbolic Logic*, 24 (1): 1-14.

–––, 1963, Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16: 83-94.

Ladyman, J., Ross, D., Spurrett, D., Collier, J., 2007, *Everything Must Go: Metaphysics Naturalized*, (Oxford: Oxford University Press).

Landauer, R., 1987, Computation: A Fundamental Physical View, *Physica Scripta*, 35: 88–95.

–––, 1991, Information Is Physical, *Physics Today*, 44: 23–29.

–––, 1996, The Physical Nature of Information, *Physics Letters A*, 217: 188.

Landauer, R. & Bennett, C. H., 1985, The Fundamental Physical Limits of Computation, *Scientific American*, July: 48–56.

Larson, A. G. & Debons, A. (eds.), 1983, *Information Science in Action: System Design. Proceedings of the Nato Advanced Study Institute on Information Science, Crete, Greece, August 111, 1978*, (The Hague: M. Nijhoff).

Laudan, L., 1981, A Confutation of Convergent Realism, *Philosophy of Science*, 48: 19–48.

Leng, M., 2010, *Mathematics and Reality*, (Oxford: Oxford University Press).

Losee, R. M., 1997, A Discipline Independent Definition of Information, *Journal of the American Society for Information Science*, 48 (3): 254-269.

Machlup, F. & Mansfield, U., (eds.), 1983, *The Study of Infomation: Interdisciplinary Messages*, (New York: Wiley).

Mayo, D., 1996, *Error and the Growth of Experimental Knowledge*, (Chicago: University of Chicago Press).

Morgan & Boumans, M. J., 2004, Secrets Hidden by Two-Dimensionality: The Economy as a Hydraulic Machine, In: S. de Chadarevian and N. Hopwood (eds.): *Model: The Third Dimension of Science*, (Stanford: Stanford University Press): 369–401.

Muller, F. A., 2010, The Characterisation of Structure: Definition versus axiomatization, in F. Stadler, D. Dieks, W. Gonzales, S. Hartmann, T. Uebel & M. Weber (eds.), *The Present Situation in the Philosophy of Science*, (Springer): 399-416.

–––, 2011, Reflections on the revolution at Stanford, *Synthese*, 183: 87-114.

Musgrave, A., 1981, Unreal Assumptions' in Economic Theory: The F-Twist Untwisted, *Kyklos*, 34: 377–387.

Nagel, E., 1961, *The Structure of Science: Problems in the Logic of Scientific Explanation*, (New York: Harcourt, Brace, and World).

Pincock, C., 2011, Modeling Reality, *Synthese*, 180: 19-32.

–––, 2012, *Mathematics and Scientific Representation*, (Oxford: Oxford University Press) (Paperback, Dec. 2014.).

Prior, A., 1971, *Objects of Thought*, (Oxford: Clarendon Press).

Quine, W.V.O., 1970, *Philosophy of Logic*, (Englewood Cliffs, NJ: Prentice Hall).

Rayo, A., 2013, *The Construction of Logical Space*, (Oxford: Oxford University Press).

Searle, J. R. 1980, Minds, Brains, and Programs, *Behavioral and Brain Sciences*, 3 (3): 417–457.

Searle, J. R. 1990, Is the Brain a Digital Computer?, *Proceedings and Addresses of the American Philosophical Association*, 64: 21–37.

Sellars, W., 1956, Empiricism and the Philosophy of Mind, in *Minnesota Studies in the Philosophy of Science*, vol. I, H. Feigl & M. Scriven (eds.), (Minneapolis, MN: University of Minnesota Press): 253–329. (Originally presented at the University of London Special Lectures in Philosophy for 1956 as "The Myth of the Given: Three Lectures on Empiricism and the Philosophy of Mind").

–––, 1962, Philosophy and the Scientific Image of Man, in *Frontiers of Science and Philosophy*, Robert Colodny (ed.) (Pittsburgh, PA: University of Pittsburgh Press, 1962): 35–78.

Schaffner, K. F., 1969, The Watson-Crick Model and Reductionism, *The British Journal for the Philosophy of Science*, 20 (4): 325–348.

Shannon, C. E., 1948, A Mathematical Theory of Communication, *Bell System Technical Journal*, 27: 379–423, 623–656.

–––, 1993, *Collected Papers*, N. J. A. Sloane and & A. D. Wyner (eds.), (Los Alamos, Ca: IEEE Computer Society Press).

Shannon, C. E. & Weaver, W., 1949, *The Mathematical Theory of Communication*, (Urbana: University of Illinois Press).

Smokler, H., 1966, Informational Content: A Problem of Definition, *The Journal of Philosophy*, 63 (8): 201–211.

Solomonoff, R.J., 1960, A preliminary report on a general theory of inductive inference, Techical Report ZTB-138, Zator.

–––, 1964a, A Formal Theory of Inductive Inference Part I, *Information and Control*, 7 (1): 1–22.

–––, 1964b, A Formal Theory of Inductive Inference Part II, *Information and Control*, 7 (2): 224–254.

–––, 1997, The Discovery of Algorithmic Probability, *Journal of Computer and System Sciences*, 55 (1): 73–88.

Spector, M., 1965, Models and Theories, *British Journal for the Philosophy of Science*, 16 (62): 121-142.

Stenius, E., 1960, *Wittgenstein's Tractatus: A Critical Exposition of its Main Lines of Thought*, (London: Blackwell).

Suárez, M., 2003, Scientific representation; against similarity and isomorphism, *International Studies in the Philosophy of Science*, 17 (3): 225-244.

——, 2004, An inferential conception of scientific representation, *Philosophy of Science*, 71 (5): 767-779.

——, 2008, *Fictions in Science: Philosophical Essays on Modeling and Idealization*, (Routledge Studies in the Philosophy of Science).

——, 2010, Scientific Representation, *Philosophy Compass*, 5 (1): 91-101.

——, 2010, Scientific Realism, the Galilean Strategy, and Representation, in in W. González (Ed.), *Scientific Realism and Democratic Society: The Philosophy of Philip Kitcher*, Poznan Studies in the Philosophy of the Sciences and the Humanities, (Amsterdam: Rodolpi).

——, 2015, Deflationary Representation, Inference, and Practice, *Studies in History and Philosophy of Science*: 4936-4947.

Suárez, M. & A. Solé, 2006, On the Analogy between Cognitive Representation and Truth', *Theoria*, 55: 39-48.

Suppes, P., 1960, A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences. *Synthese*, 12 (2-3): 287-301.

——, 1970, Probabilistic Grammars for Natural Languages. *Synthese*, 22 (1-2): 95 - 116.

——, 1988, Representation Theory and the Analysis of Structure. *Philosophia Naturalis*, 25 (3/4): 254.

——, 2002, Representation and Invariance of Scientific Structures. Models and Idealization in *General Philosophy of Science*.

Suppes, P. & Tarski, A., 1962, *Logic, Methodology and Philosophy of Science Proceedings of the 1960 International Congress* (Stanford University Press).

Thompson, E., 2007, The enactive approach, *Mind in life* (Paperback ed.), (Cam., Mass: Harvard University Press).

Toon, A., 2010, The Ontology of Theoretical Modelling: Models as Make-Believe, *Synthese*, 172: 301–315.

van Benthem and Martinez, 2008, Stories of Logic and Information, in Adriaans and van Benthem, 2008b.

Vedral, V., 2010, *Decoding Reality: The Universe As Quantum Information*, (Oxford University Press: Oxford).

Vereshchagin, N.K. & Vitányi, P.M.B., 2004, Kolmogorov's Structure functions and model selection, *IEEE Transactions on Information Theory*, 50 (12): 3265–3290.

Warner, T., 1996, *Communication Skills for Information Systems*, (London: Pitman Publishing).

Wiener, N., 1954, *The Human Use of Human Beings: Cybernetics and Society*, (Boston: Houghton Mifflin); reissued in 1989 with a new introduction by S. J. Heims, (London: Free Association).

Weisberg, M., 2012, Getting Serious About Similarity, *Philosophy of Science*, 79 (5):785-794.

Wheeler, J. A., 1990, Information, Physics, Quantum: The Search for Links, in *Complexity, Entropy, and the Physics of Information*, W. H. Zureck (eds.), (Redwood City, CA: Addison Wesley).

Wittgenstein, L., 1922, *Tractatus Logico-Philosophicus (TLP)* C. K. Ogden (trans.), (London: Routledge & Kegan Paul). Originally published as "Logisch-Philosophische Abhandlung", in *Annalen der Naturphilosophische*, XIV (3/4), 1921.

–––, 2009, *Philosophical Investigations (PI)*, 4th edition, P.M.S. Hacker and J. Schulte (eds. and trans.), (Oxford: Wiley-Blackwell).