

Intentional Models as Essential Scientific Tools

Eric Hochstein

Abstract: In this paper, I argue that the use of scientific models which attribute intentional content to complex systems bears a striking similarity to the way in which statistical descriptions are used. To demonstrate this, I compare and contrast an intentional model with a statistical model, and argue that key similarities between the two give us compelling reasons to consider both as a type of phenomenological model. I then demonstrate how intentional descriptions play an important role in scientific methodology as a type of phenomenal model, and argue that this makes them as essential as any other model of this type.

1. Introduction

Intentional concepts (e.g. “representations”, “goals”, “beliefs”, “information”) are commonly found as part of scientific models in many different domains. Philosophers of mind and philosophers of science, however, have long debated the scientific merits of such models. These debates often hinge, either explicitly or implicitly, on whether we can provide a suitable metaphysical account of the phenomenon of intentionality, and whether the intentionality ascribed to a system is real. Only under these conditions is the application of an intentional model thought to be scientifically appropriate (see, for instance: Fodor 1975, 1987, 2008; Field 1978; Searle 1980, 1992; Dretske 1981; Stich 1983; Churchland 1981; Sarkar 1996, 2000, 2004; Adams 2003; Floridi 2005; Weber 2005).

While some philosophers, like Daniel Dennett, argue for a more instrumental account of intentional descriptions in science, such accounts often propose that intentional ascriptions have, at best, a very limited role to play in our scientific practices. Dennett, for instance, claims that while intentional characterizations are useful for generating predictions in day-to-day life, they only function as a heuristic device in scientific contexts; one that we ought not take “*too seriously*” (Dennett 1987, p. 350, emphasis in text). Given this, he insists that intentional characterizations of human behaviour are ultimately “*vacuous as psychology*” (1971, p.99).¹

In this paper, I argue that the use of scientific models which attribute intentional content to systems (hereafter “*intentional models*”) are critical to our scientific understanding of systems, irrespective of whether we believe the system has genuine metaphysical intentionality or not. Instead, I propose that such intentional models are important because they are used analogously to the way in which we use phenomenological models.² Given this, I propose we have good reason to view intentional models *as* a species of phenomenological model.

More specifically, I propose that if we look at the way in which intentional models are used within scientific domains, we find that they share a great deal in common with the way in which we use statistical models. Given similarities between the two, I propose that we have compelling evidence that intentional models are best interpreted as a particular type of phenomenological model, just as statistical models are. When understood in this way, the scientific value of intentional models can be made clear. Intentional models play an important

¹ It should be noted that, given the prolific nature of Dennett’s work, he is not always consistent regarding this point. Overall, however, the idea that intentional descriptions play a restricted role in science has been a strong undercurrent throughout Dennett’s work.

² Given that the focus of this paper is specifically on the scientific benefits of intentional concepts as part of scientific models, I remain agnostic as to whether the everyday use of intentional language in colloquial contexts can be characterized as a phenomenological model as well. There is, however, recent evidence that even the folk psychological use of intentional idioms can be understood as something akin to the application of a model, albeit not necessarily a scientific one (Maibom 2003, 2009; Godfrey-Smith 2004, 2005). If such accounts are true, then the position I advocate here may apply to intentional descriptions more broadly.

role in our discovery and understanding of the implementation of complex systems, and are needed to identify features of systems that detailed structural or mechanistic models cannot identify. Crucially, the successes of such models are not dependent on whether one is a realist, or an instrumentalist, regarding intentionality or intentional states. The value of intentional models does not stem from their metaphysical insights, but their epistemic role in scientific practice as a type of phenomenological model.

I begin in Section 2 by looking at the various sorts of intentional models that exist in different scientific domains. Then, in section 3, I examine the nature of phenomenological models, and explain why statistical models are phenomenological models. In Section 4, I select an intentional model and a statistical model to compare and contrast. In Section 5, I highlight four key similarities that exist between the two types of models, and show why intentional models are importantly similar to statistical models to warrant being considered a type of phenomenological model. I then demonstrate how intentional models are essential to science as a phenomenological model. Lastly, in Section 6, I look at some possible differences between the two types of models that may threaten this idea. I demonstrate that these differences are too shallow to challenge the account on offer.

2. Intentional Models in Science

While intentional models are most commonly associated with psychology (often used as a means of explaining human behaviour via the attribution of mental representations such as “beliefs”, “intentions” and “desires”), they are hardly limited to this domain. Intentional models of various sorts can be found in a wide range of scientific domains, from psychology (Ajzen 1985, 1988,

1991), to neuroscience (Harris, Sheth, & Cohen 2008; Andersen & Cui 2009; Eliasmith et al. 2012), to artificial intelligence (Parisien & Thagard 2008), to biology (Sterelny, Smith, & Dickison 1996; Sterelny & Griffiths 1999; Maynard Smith 2000), and even to virology:

We sometimes speak, in a “loose and popular” way, as though viruses have human-like beliefs and desires. [...] It comes as no surprise when such intentional idioms occur in the hyperbolic prose of the mainstream media. However, it is more surprising when scientific discourse relies on anthropomorphizing claims when discussing the nature of viruses. In recent papers, it has been suggested that if the HIV “virus wants to resist AZT, it needs to make a specific mutation at codon 215 and at another position, such as position 41” and that the Hepatitis C virus must try to coexist with its host by reducing its visibility. (Huebner 2011, p. 441)

For another example, consider the use of intentional models in neuroscience. In the neuroscience of vision, scientists often characterize various neurological mechanisms in terms of the intentional information they contain. Certain neurons within a frog’s visual system, for instance, are thought to *represent*, or contain intentional *information about*, darting black objects in their environment (Lettvin et al. 1959). Such appeals to intentional terminology (“representations”, “information”) are commonplace in neuroscience.

According to many philosophers, the question of whether or not such intentional models have an important role to play within scientific practice depends on whether we can tell an appropriate metaphysical story about intentionality. Take the case of intentional models in virology. According to Huebner:

Philosophical reflection suggests that although we can use intentional idioms whenever we find an entity that exhibits robust and systematic patterns of behavior, a plausible theory of mentality should not treat viruses as even minimally minded. After all, viruses are nothing more than packets of RNA (or sometimes DNA) encased in protein. [...]

Viruses are not the kind of entities that can want to threaten our wellbeing; they do not make plans and adopt deceptive strategies for navigating our immune systems; and they have not learned to outsmart our best anti-virals. Although many viruses exhibit systematic and predictable patterns of behavior, we can and should treat them as complex biological robots. (2011, p. 441)

In this case, Huebner proposes that it is inappropriate to apply intentional models to virology because *viruses do not have intentional states*. Next, consider models in molecular biology that attribute intentional information to genes. According to Marcel Weber:

I want to exhibit the deeper metaphysical reasons why some common ways of describing the causal role of genes in development and evolution are problematic. Specifically, I show why using the concept of information in an intentional sense in genetics is inappropriate, even given a naturalistic account of intentionality. Furthermore, I argue that descriptions that use notions such as programming, directing or orchestrating are problematic not for empirical reasons, but because they are not strictly causal. They are intentional. (2005, p.407)

As with Huebner above, Weber argues that we ought not use intentional models in genetics because we have no reason to assume genes possess the appropriate *metaphysical* properties required for intentionality. A similar argument is also made by Sarkar (2004, p.262). In other words, the only way to justify the use of an intentional model in a given scientific domain is by first providing a metaphysical account of intentionality. Only then we can determine what things genuinely have it, and whether we are justified in describing them that way. Yet, I propose that this way of thinking about intentional models is more of an obstacle to scientific progress than a part of it. If biologists can, and do, use intentional models, then the question is whether such models are useful, and why? The issue of metaphysics is another matter altogether. The value of intentional models, I propose, is that they behave within scientific practice as a type of phenomenological model.

3. Phenomenological Models and Their Applications

In order to demonstrate that intentional models behave akin to phenomenological models, I begin this section by providing a brief explanation of what phenomenological models are, and how they are typically used in scientific contexts. I then look at two different types of phenomenological models to further highlight their similar scientific usage. This will help set the stage for a comparison between an intentional model and a standard phenomenological model in the following sections.

3.1 Phenomenological models

Phenomenological models are defined by their ability to characterize, and predict, a particular phenomenon (such as the behaviour of a complex system) without attempting to decompose it into physical parts and operations for better understanding. According to Frigg and Hartmann (2009), “a traditional definition takes [phenomenological models] to be models that only represent observable properties of their targets and refrain from postulating hidden mechanisms and the like.”³ A more detailed description, offered by Carl Craver, is that:

...all one requires of a [phenomenological] model is that it be *phenomenally adequate*. That is, the input–output mapping in S [e.g., a postulated algorithm, function, or account, that generates a mapping from inputs to outputs] should be sufficiently similar to the input–output mapping in T [e.g., the observed regularity in the actual input-output of the system] for one’s needs. Few models are actually isomorphic with the phenomenon, given that models typically abstract away from the precise details of the system being modeled, that they typically are only approximate, and that they make simplifying assumptions in order to apply a particular formalism. (2006, p. 357)

Many consider such models to be largely unexplanatory in the sciences of the mind due to their focus on characterizing and predicting phenomena instead of developing a *mechanistic*

³ It should be noted that this particular definition is vague. It could be interpreted as saying that a phenomenological model is used as a means of characterizing directly observable feature(s) of a system, such as behavioural regularities, while saying nothing about the unobserved causes of those features/regularities. On the other hand, it could be interpreted as saying that a phenomenological model is one that only invokes entities and relations that are directly observable (while refraining from references to theoretical posits, or inferred entities). I propose that the former definition is more in-line with actual scientific usage than the latter. Consider that statistical models are commonly considered to be phenomenological models, yet they posit boundary conditions which are not directly observable features of systems. Likewise, many examples of phenomenological models knowingly invoke theoretical unobserved entities. Carl Craver (2006, p.356), for instance, argues that one can use a Ptolemaic model of the solar system as a phenomenological model if one is only interested in a predicatively adequate account of where the planets will appear in the sky. Yet such a model posits numerous unobservable theoretical entities such as equants, deferents, and epicycles (ibid, p. 358). What makes such a model phenomenological is that it is used as a means of describing aspects of the systems that are directly observable (like where the planets will appear in the sky) while saying nothing about underlying causes, not that it only makes reference to observable entities. Many scientific examples of phenomenological models are of this sort.

explanation of their production (see: Machamer et al. 2000; Bunge 2003; Bechtel 2008; Darden 2006; Craver 2006; Wimsatt 2007). Craver, for example, claims that:

A model can be richly phenomenally adequate and non-explanatory. This is the take-home lesson of the several decades of attacks on covering-law models of explanation at the hands of advocates of causal–mechanical models of explanation: merely subsuming a phenomenon under a set of generalizations or an abstract model does not suffice to explain it. (2006, p.357-359)

It should be noted that this view, while widely held, is not without detractors. The question of whether phenomenological models can be explanatory in the behavioural sciences is still in debate. Some, for instance, have proposed that we need a more pluralistic understanding of explanation in order to better account for actual scientific practice in the behavioural sciences (Longino 2006; Chemero & Silberstein 2008). What is important to note for the purposes of this paper, however, is the fact that we use phenomenological models for different scientific purposes than those we use mechanistic models for, irrespective of whether each can be explanatory under the appropriate conditions.

While mechanistic models characterize the physical objects and operations that make up a system, phenomenological models are used instead to provide detailed accounts of the phenomena produced by mechanisms, to summarize data, to identify similarities that exist across vastly different types of mechanistic systems, to identify behavioural regularities, and to make predictions (Batterman 2002; Bogen 2005; Craver 2006).⁴

⁴ A note of clarification: The way in which I describe a mechanistic model suggests that I am equating “mechanisms” with physical parts and operations. Yet, it might be argued that many appropriate scientific uses of the term “mechanism” in the life sciences need not be characterized this way. Within psychology, for example, one might make reference to learning mechanisms without providing any story of how biological or neurological objects are interacting. These are certainly legitimate uses of the term “mechanism”, and my intention is not to stipulate

3.2 The Different kinds of phenomenological models

Phenomenological models come in all different shapes and sizes, each useful for characterizing and predicting different sorts of phenomena. Statistical models and some dynamical models, for instance, provide examples of phenomenological models. Dynamical models involve the application of Dynamic Systems Theory (DST) to predict the behaviour of complex systems as they change through time. These models use differential equations in order to represent the system as a vector moving through a multi-dimensional state space. Dynamical models employed by those who embrace a specific form of DST in cognitive science (Busemeyer & Townsend 1993; Thelen & Smith 1994; Van Gelder & Port 1995) typically use such models phenomenologically: it is “usually only observable behaviour [that] is mapped to the model” (Eliasmith, 2010, p. 319). In this respect, many dynamical models in cognitive science act as

correct usage or chastise other common uses of the term. Instead, my intention is highlight the fact that certain types of scientific models (those commonly referred to as “mechanistic” models) are adopted for very particular purposes: to identify the physical implementation of a given system. As Eliasmith notes:

In the case of cognitive and brain sciences, useful explanations are those that appeal to subpersonal mechanisms [understood in terms of physical parts and operations]. This is because it is precisely such explanations which provide a basis for both intervention in behaviour and the artificial reproduction of those behaviours. These mechanisms must be specific enough to allow for intervention. That is, the mechanisms must be specified in a way that relates to the measurable and manipulable properties of the system. (2010, pp. 316)

If we are only talking about a mechanism at an abstract level without appealing to parts and operations, our model may tell us behavioural patterns and regularities, but provides no insight into why such regularities exist as they do or how they might change under different conditions. For this reason, some argue that scientific models which employ the more abstract usage of the term “mechanism” tend to behave more as phenomenological models (Machamer et al. 2000; Craver 2006). In contrast, the application of mechanistic models are expressly employed for the purposes of telling us how a given system is being physically instantiated, and thus provides an account of *why* the system behaves as it does. It is this usage of the term “mechanism”, and “mechanistic model”, that I am appealing to for the purposes of this paper (for more details, see: Machamer et al. 2000; Bunge 2003; Bechtel 2005, 2007, 2008; Bechtel and Abrahamsen 2007; Glennan 2005; Darden 2006; Craver 2006, 2007; Thagard 2006, 2009, 2012; Wimsatt 2007; Eliasmith 2010).

phenomenological models, providing no mechanistic explanation for the behaviour of the system.⁵

Statistical models, meanwhile, act as a different sort of phenomenological model. While dynamical models generate predictions by applying an appropriate set of differential equations, statistical models typically predict by assigning various probabilities to measurable states of the system given known conditions. This allows us to identify regularities in the data we collect about systems, and to use them as a means of generating predictions (for details, see Eliasmith 2010).

Given that different sorts of phenomenological models characterize systems in different ways, one type of model may be more useful than another for our purposes given the particular dataset we have available to us at any given time. We might, for instance, lack the relevant information needed to construct a dynamical model of a system, but not a statistical one. Or vice versa.

Statistical models and dynamical models of the sort discussed above are united by the fact that they stay silent in regards to the physical implementation of the systems they describe, and are instead used for tasks such as prediction, data analysis, and the identification of behavioural patterns. If I am to show that scientific ascriptions of intentional content function as a type of phenomenological model, then I need to demonstrate how intentional models are crucially similar to dynamical and statistical models to warrant being subsumed under the same category of model. However, I must also show how they differ so as to make intentional models useful in scientific contexts where dynamical and statistical models are not. To accomplish this, I begin by selecting an intentional model, and a statistical model, to compare and contrast.

⁵ Given that we use dynamical models for different purposes than those we use mechanistic models for, there is often thought to be no tension between the two types of models. For more details, see Kaplan & Bechtel 2011.

4. Statistical and intentional models

To argue that statistical and intentional models share crucial similarities, in this section I briefly introduce an example of each kind of model. This introduction is followed by an examination of the similarities between these examples, where additional details about each model are provided.

4.1 Example of a Statistical Model

As an example of a statistical model, consider what is often called the “Thompson Effect”. This is the phenomenon in which objects appear to us to move faster or slower based on how greatly the object contrasts with its surrounding environment. Alan Stocker and Eero Simoncelli (2006) developed a statistical model that can accurately predict how fast an object will appear to be moving to an observer given the level of contrast between the object and its surrounding environment. They did this by altering the modern framework of statistical estimation used to traditionally model perception. This traditional framework viewed an optimal observer in terms of two probability distributions:

First, the variability of a set of measurements, m , is specified as a conditional probability distribution, $p(m/v)$, where v is the stimulus speed. The variability is due to a combination of external sources (e.g., photon noise) as well as internal sources (e.g., neural response variability). When considered as a function of v for a particular measurement, this conditional density is known as a likelihood function. The second component is a prior probability distribution, $p(v)$, which specifies the probability of encountering stimuli moving at any particular retinal speed. According to Bayes’ rule, the product of these two

components (when appropriately normalized) gives the posterior distribution. Common choices are the mean, or the mode. Biases in the perceived speed of low-contrast moving patterns arise intrinsically with this model, assuming a prior that favors low speeds: Lower contrast stimuli lead to noisier measurements, producing a broader likelihood function, which leads to a lower speed estimate. (Stocker & Simoncelli 2006, p. 578)

The problem with this traditional model is that it is extremely difficult to experimentally determine what the prior distribution and the likelihood function actually are. Similarly, there are constraints on the model that cast doubt on its ability to accurately represent the phenomenon in question (Stocker & Simoncelli 2006, p. 578). To compensate for this, Stocker and Simoncelli embed a Bayesian estimator into the traditional model. This estimator is calibrated based on the trial-by-trial responses of subjects observed in a forced choice speed discrimination experiment. According to Stocker & Simoncelli:

We were able to validate the ability of a Bayesian observer model to account for the data, and also to determine the prior distribution and internal noise level associated with the best-fitting Bayesian estimation model. (2006, p. 579)

By embedding the Bayesian estimator into the traditional model of perception, Stocker and Simoncelli were able to generate a predictive model that overcomes the shortcomings of previous statistical accounts of the phenomenon. Most importantly for our purposes, however, is the fact that their model does not provide any mechanistic story for what generates the Thompson Effect, but still accurately predicts the performance of subjects in a wide range of tasks. In this regard, statistical models like Stocker & Simoncelli's model function as a type of phenomenological model.

4.2 Example of an Intentional Model

For an example of an intentional model, let us turn to molecular biology. It is common practice in molecular biology to attribute information to genes regarding phenotypic traits. However, the way in which “information” is understood can often be unclear. There are technical, non-intentional, senses of information associated with communication theory (such as Shannon-Weaver information), and models that attribute this sort of information to genes do have an important place in biology. However, the models that attribute phenotypic information to genes are often not usefully characterized in terms of Shannon-Weaver information (Sterelny & Griffiths 1999, p. 104; Godfrey-Smith 2000; Maynard Smith 2000; Sterelny 2000, p. 196; Sarkar 2004; Weber 2005). Instead, these biological models are best understood as employing a distinctly *intentional* notion of information:

A single nucleotide substitution can destroy the genetic information of a gene, for example, if it leads to the insertion of a stop codon. But the mutant DNA sequence still contains the same amount of Shannon-Weaver information. What this suggests is that ‘information’ is actually used in an *intentional* or *semantic* sense, that is, in a similar or even the same sense in which we say about English sentences that they contain information. (Weber 2005, p.410).

While Weber ultimately calls into question the scientific merits of these models given that we have no *metaphysical* grounds to suppose that genes have genuine semantic properties, many

have espoused the essential *methodological* virtues of employing such intentional models in biology, irrespective of their metaphysical standings (Sereno 1991; Sterelny, Smith, & Dickison 1996; Ji 1997, 1999; Sterelny & Griffiths 1999; Maynard Smith 2000; Sterelny 2000).⁶ I will demonstrate that the virtues associated with such models mirror the virtues that apply to statistical models like Stocker & Simoncelli's model. Moreover, the metaphysical question of whether genes genuinely have intentionality is ultimately irrelevant to whether the model has the scientific virtues of a phenomenological model.

5. Similarities Between Intentional and Statistical Models

In this section I argue for four crucial traits that are shared by intentional and statistical models. I will use the Stocker and Simoncelli statistical model, and information models from molecular biology, to demonstrate how both types of models exhibit these traits.

5.1 Neither model directly describes the physical mechanisms of a system

Statistical models, being a species of phenomenological model, characterize systems without telling us about their underlying causal mechanisms. As Eliasmith notes:

⁶ It is also important to note that the attribution of intentional information to genes is not a case of mere casual anthropomorphism on the part of biologists, but are part of genuine scientific models. As John Maynard Smith notes:

Transcription, translation, code, redundancy, synonymous, messenger, editing, proofreading, library - these are all technical terms in biology. I am not aware of any confusions arising because their meanings are not understood. In fact, the similarities between their meanings when referring to human communication and genetics are surprisingly close. (Maynard Smith 2000, p.178)

Statistical descriptions are highly implementation independent. Statistical models focus on describing the regularities in the data and hence are silent with respect to the particular physical implementation. In essence, these descriptions would not change if the implementation changes and statistical properties do not. (2010, p. 315)

Intentional models are similarly implementation independent. To say that the haemoglobin-S gene carries information, or instructions, for the sickle cell trait does not, by itself, tell us what mechanisms are at work in the implementation of these “instructions”. This is why John Maynard Smith points out that “the weakness of these models [...] is that they do not tell us where the ‘rules’ [for these genetic programs or instructions] come from” (2000, p.190). They do, however, allow us to identify regularities in the data we gather about biological system when we know relevant environmental conditions.

5.2 Both models generate predictions of systems despite a lack of mechanistic assumptions

We use both statistical and intentional models to form predictions of systems whose mechanisms we do not yet understand, or cannot identify. Stocker and Simoncelli’s statistical model can predict the Thompson Effect despite a lack of information regarding the actual mechanisms generating it. In fact, one of the primary *benefits* of statistical models is that they allow us to generate predictions of systems with unknown mechanisms. We often explicitly use statistical models for the purpose of modeling systems that we cannot explain or predict mechanistically (Eliasmith 2010, p.315).

The same is true of intentional models. The use of intentional models in genetics are often used to make powerful and important predictions. Sereno, for instance, notes that the use of

intentional models in genetics can be used “to make preliminary predictions about [...] some presently hidden mechanisms” (1991, p.19). Maynard Smith similarly argues that it is common practice for biologists to attribute intentional information (in the form of the genetic code) to DNA as a means of predicting amino acid sequences (2000, p.184). Likewise, Godfrey-Smith points out that through the application of such models, “generalizations can be made about the division of labor between the different kinds of maco-molecules in cells (proteins, carbohydrates, lipids, and nucleic acides)” (2000, p.204). In these cases, even though the intentional model by itself does not describe the physical implementations of the system, it allows us to generate relevant scientific predictions.

5.3 Both models are used to help us learn about unknown causal mechanisms

If we are trying to learn about the unknown mechanisms of a system, one good way to do this is by understanding the behavioural regularities produced by these mechanisms. Statistical and intentional models are useful for exactly this purpose. Our ability to effectively characterize and predict the behaviour of a system allows us to narrow the list of possible mechanisms that might be implementing it. Take, for example, Stocker & Simoncelli’s model. While their account does not tell us what the mechanisms underlying the system are, it does provide insights into possible mechanisms:

The form of the contrast-dependent measurement noise in our model suggests that the locus of representation for measurements m is likely to be cortical. Neurons in area MT are a natural choice: They are highly motion selective, and their responses have been directly linked to perception. (2006, p. 583)

Even though Stocker & Simoncelli's model does not directly describe the neurological mechanisms responsible for the Thompson Effect, by describing and predicting the phenomenon in a detailed way, it provides insights into what the mechanisms must be like. The more detailed our account of the phenomenon, the more it tells us about what the unknown mechanisms must be capable of producing given known constraints. And this in turn helps us narrow the field of possible mechanisms. In this case, it means that whatever mechanisms are responsible for producing the Thompson Effect must meet the regularities observed by Stocker & Simoncelli's model.

Now let us turn to intentional models. Historically, intentional models that characterized genes in terms of a genetic code were pivotal in the discovery of causal mechanisms in molecular biology. For example:

[The] discovery of the relationship between DNA and protein –as a triplet code in which the correct 'reading frame' is maintained by accurately counting off in threes, and whose meaning can be destroyed by a 'frame shift' mutation [...] arose from the coding analogy.
(Maynard Smith 2000, p. 184)

Even those who are skeptical of the role that such intentional models currently play in biology are quick to acknowledge that such models have proven extremely important in the development of molecular biology as a means of discovering and understanding the mechanistic underpinnings of biological systems (see, for instance, Sarker 2000, 2004; Weber 2005).

Intentional models, in virtue of being predictive, provide us with essential information that allows us to narrow down the set of possible mechanisms producing those behaviours. Given environmental restrictions, not any implementation of a physical system will be capable of producing the behaviours identified and predicted by an intentional model under those conditions. In this respect, by highlighting behavioural regularities displayed by the system given known environmental conditions, intentional models help us to reverse engineer systems by identifying constraints on the set of possible causal mechanisms.

5.4 Generating a mechanistic account of a system does not necessarily make either model obsolete

While both statistical and intentional models are used to help generate mechanistic accounts of systems, this does not imply that such models become obsolete the moment we have a mechanistic account in hand. We use different scientific models for different purposes. And while mechanistic models are ideal for characterizing the structure of systems, this does not mean they are ideal for other sorts of purposes.

For example, the model that is best for characterizing the mechanisms of a system is not necessarily the best model to use for generating predictions. As Mark Wilson notes, “the reasoning requirements natural to design tasks are often quite different than those pertinent to prediction et al. and greatly influence the descriptive vocabulary we find suitable” (Wilson, 2006, p. 326). Knowing the mechanisms that produce the Thompson Effect, for instance, does not mean that we do not use a statistical model like Stocker & Simoncelli’s in order to predict the effect. Practical constraints on scientific practice often mean that statistical models are ideal for

scientific predictions when restrictions on time and resources make mechanistic models inefficient or impractical. And the same lesson applies to intentional models.

Consider the problems that faced the neuroscience of vision in the mid-20th century. The neurological mechanisms responsible for vision were exceedingly difficult to predict based only on a mechanistic understanding of the system. It was not until neuroscientists began attributing *intentional content* to neurological systems that relevant predictions could be made (Marr 1982; Betchel 2007, p. 184). Certain neurons, for example, were found to contain visual information about the edges of objects in one's visual field. According to David Marr, this information needed to be "analyzed and understood in a way that [was] independent of the particular mechanisms and structures that implement them in our head" (Marr 1982, p. 19). Only by intentionally modeling the system could we generate relevant predictions.

Likewise, there is evidence that intentional models in biology may possess similar virtues.⁷ Even if we could, in principle, do away with intentional models in favour of mechanistic ones in molecular biology, the pragmatic benefits intentional models provide *in practice* may justify their continued usage. Consider Sarkar's observation that:

Routinely, talk of information is intertwined with linguistic metaphors, from both natural and artificial languages. [...] The use of such talk is so pervasive that it almost seems impossible that, short of pathological convolution, the experimental results of genetics can even be communicated without these resources. (2004, p. 266)

Further evidence for this may be seen in Dennett's claim that:

⁷ Although for a contrasting view, see Sarkar 1996, 2000; Weber 2005.

Pending completion of our mechanical knowledge, we need the intentional characterizations of biology to keep track of what we are trying to explain, and even after we have all our mechanical explanations in place, we will continue to need the intentional level against which to measure the bargains Mother Nature has struck. (1987, p.315)

In a similar vein, it has recently been argued that models in physics which disregard the structural details of physical systems (just as statistical and intentional models do) can have profound virtues over their mechanistic counterparts. Robert Batterman, for example, argues that a model that *ignores* many of the correct structural details of a system can often be a better candidate for characterizing certain behavioural regularities than those that provide a fine-grained account of the system's mechanisms. As an example, Batterman points to our use of the fundamental theories of statistical mechanics to characterize phase transitions, like water turning from liquid to ice:

From the point of view of the underlying fundamental theory whose proper focus is on the interactions of a *finite* number of molecular components of the macrosystems, these qualitative changes are genuinely novel. The upshot is that the statistical mechanics of finite systems is explanatorily insufficient. While it gets the ontology of blobs of gases and fluids right (they are composed of a finite number of interacting molecules), there remain macroscopic phenomena —universal patterns of behavior— that cannot be explained by this fundamental theory. (2011, pp. 1033-1034)

Put simply, a model that provides a detailed structural account of the system (identifying the finite molecules that make up the system, and their interactions) fails to identify or account for the phenomenon of phase transitions. The only way for us to successfully model this behavioural pattern of water is by interpreting the system as having an infinite number of molecules (which

we know it does not have). In other words, many of the correct structural details of the system must be left aside in order for us to accurately predict the behavioural regularity produced by the system. In such cases, mechanistic details “actually detract from an understanding of the phenomenon” (Batterman 2002, p.22).

Models like intentional and statistical models are often used to represent and predict the behavioural regularities produced by systems, while mechanistic models are used to represent their internal parts and operations. In this regard, the different types of models are used to represent different aspects of systems for different purposes, and so have different pragmatic virtues.

5.5 Given these shared traits, both models function as a species of phenomenological model

The similarities listed above between statistical and intentional models provide evidence for my contention that intentional models function within scientific practice as a type of phenomenological model. Consider that a benefit of both statistical and intentional models is that they allow us to generate predictions of systems whose mechanisms we cannot identify or do not fully understand. In other words, both models are *implementation independent*. This is one of the defining features of phenomenological models: “[They] are a means for extracting *stable phenomenologies* from unknown, and perhaps unknowable detailed theories” (Batterman 2002, p. 35, emphasis in text).

Recall that phenomenological models come in all shapes and sizes, each useful in different contexts and for different purposes (e.g. dynamical versus statistical descriptions). Intentional models are simply another breed of phenomenological model, one useful for

characterizing and predicting behaviour that other models are not well suited for. For instance, intentional models allow us to generate predictions when we do not have the luxury or ability to quantify over relevant features of the system needed to generate statistical or dynamical models.

If we understand intentional models not as descriptions of the structure of systems, but instead as a species of phenomenological model, then their role in science becomes much clearer. They play an important role in our discovery, and understanding, of the mechanisms that constitute complex systems. They similarly allow us to identify behavioural patterns, and generate essential predictions, in contexts where other sorts of models are either unavailable or unsuitable. These scientific virtues of intentional models do not in any way require that we have a metaphysical story of what intentionality is, how it works, or whether we are realists/antirealists regarding intentional states. Yet these virtues are more than enough to validate the presence of such models in scientific domains.

6. Differences between Statistical Models and Intentional Models

Up until this point, I have focused only on the *similarities* between intentional and statistical models to make my case. In order to defend the claim that they are relevantly similar, I must also rule out any key dissimilarities between the two that might disqualify intentional models from being a type of phenomenological model. In this section, I examine such dissimilarities and demonstrate that they are not a threat to the account I provide above.

6.1 We use intentional models to explain behaviour, not just predict it

The first important difference to note is that while we commonly explain behaviour by appealing to intentional models, we typically do not explain behaviour by appealing to statistical models. Stocker & Simoncelli's model is not considered to be an *explanation* of the Thompson Effect, but only a means of predicting it. Intentional models, on the other hand, are commonly used for predictions *and* explanations.

Consider intentional models in psychology. We do not only *predict* people's behaviour with the attribution of beliefs and desires, but *explain* it as well. We can say that John votes for the New Democratic Party of Canada *because* of his beliefs, intentions, and desires. In this regard, intentional models do not seem to function the way statistical models do given that they are often explanatory as well as predictive.

This distinction between statistical and intentional models is ultimately not a sufficient reason to think that intentional models are not phenomenological models. While it is true that we often do not use statistical models to generate explanations in the behavioural sciences, this does not mean that we do not use phenomenological models *of any sort* to provide explanations in these domains. It has been argued that, depending on the question being asked, and the particular phenomenon under investigation, phenomenological models often do provide us with scientific explanations (Chemero & Silberstein 2008). Similarly, Batterman argues that since some phenomena can only be seen and characterized using phenomenological models, that we ought to consider such models to be explanatory despite their mischaracterizations, or agnosticism, regarding the underlying implementation of the system (Batterman 2000, 2002).

But even if we were to suppose that mechanistic models are the best explanatory accounts in the life sciences (as some do), this would still not be enough to challenge the idea that intentional models behave as phenomenological models in these domains. Even if

phenomenological models are insufficient to function as rigorous scientific explanations, this does not mean that they would have no explanatory power *whatsoever*. In virtue of being predictive, such models still provide a degree of explanatory power (Dennett 1987, p. 259).

According to Craver, one of the reasons that mechanistic models are explanatory in the life sciences is that they “afford the ability to say not merely how the system in fact behaves, but to say how it *will* behave under a variety of interventions” (Craver 2006, p. 358, emphasis in text). Explanations tell us not only what the system’s actual behaviour is, but also what it would be in counter-factual situations. Phenomenological models, in virtue of summarizing and predicting the phenomenon under investigation, still “typically allow one to answer some [“what-if-things-had-been-different”]-questions” (Craver, 2006, p. 358). In this respect, such models still provide us with limited explanations, even if we work on the assumption that they are always insufficient to provide many robust scientific explanations. And so the fact that intentional accounts are used to provide explanations does not disqualify them as phenomenological models, it simply means that they might provide only a constrained or limited sort of explanation.

6.2 Intentional models are normative, statistical models are not

It has been argued that intentional models are ill-equipped to play a role in scientific practice because they make predictions based not on how systems genuinely function, but only on how the system *ought* to behave *if* it were rational. In this respect, unlike statistical models, intentional models are inherently *normative*, and not descriptive like other models in science. To

better understand this particular objection, let us consider intentional models from domains like psychology:

There is a normative character to our practices of mental state ascription that is foreign to the theories involving natural kinds in the sciences. The normative aspect of mental state ascription can be seen in several related ways. Most generally, when we ascribe mental states, we do so against the background assumption that we are dealing with a *rational* agent; i.e., we attribute propositional attitudes to an agent against our background conception of what she *ought* to believe and desire. [...] Attributing a completely irrational set of beliefs to an agent defeats the purpose of belief attribution, and the attribution itself loses sense. [...] However, scientific theorizing does not appeal to overtly normative standards in the way that mental state ascription does. Our choice of theories is not guided by an ideal conception of how the world ought to behave. (Sehon 1997, p.334).

The problem with this sort of objection is that it mischaracterizes actual scientific practice. I propose that intentional models are in fact no different from our use of other kinds of idealized models in science. Recall the example of phase transitions discussed in section 5.4. When we use statistical mechanics to model the behaviour of phase transitions, we *are* guided by an ideal conception of how the world ought to behave. Instead of modeling the system the way it genuinely is, we model it against the background assumption that it can have properties like an infinite number of molecular components, or infinite volume. So just as intentional models only describe the way a system *ought* to act *if* it were rational, so too does statistical mechanics only

describe how a system *ought* to act *if* it had an infinite volume. In neither case is the model strictly descriptive, since idealizations are required in both cases to make the relevant scientific predictions. And so I propose that the “rationality” case involving intentional models is just another variety of the same sort of idealization found in phenomenological models in other domains of science.

In summation, the potential dissimilarities discussed in this section fail to provide compelling evidence for the idea that statistical models and intentional models cannot be examples of the same type of scientific model.

7. Conclusion

I have argued that intentional models are best understood as a type of phenomenological model when used in scientific practice. To provide support for this idea, I have highlighted the numerous similarities that exist between intentional models and statistical models. Both models are implementation independent. Both models are used to generate predictions of complex systems with unknown mechanisms. Both models play an important role in the discovery and understanding of the mechanisms that make up complex systems. And both models are used in contexts where mechanistic models are unhelpful or unavailable. These are all key features of phenomenological models. Likewise, these virtues do not require that we take any stand whatsoever regarding the realism of intentionality or intentional states.

This means that, regardless of whether we wish to be realists or anti-realists about intentionality, intentional *models* play an essential practical role in our scientific practices. The

information they provide is invaluable in helping us to learn how complex systems are implemented, and how they behave.

Acknowledgements:

I am very grateful to Chris Eliasmith and Doreen Fraser for feedback and helpful discussions on early drafts of this manuscript. An earlier version of this paper was also presented at the 2012 Annual Congress of the Canadian Philosophical Association, where audience questions and comments were particularly insightful and helpful in shaping the paper. Lastly, I would like to thank the three anonymous referees of this journal whose comments and feedback helped to strengthen and improve the paper significantly.

References

- Adams, F. (2003) The Informational Turn in Philosophy. *Minds and Machines* 13: 471-501.
- Andersen, R & Cui, H. (2009). Intention, Action Planning, and Decision Making in Parietal-Frontal Circuits. *Neuron* 63 (5): 568-583.
- Ajzen, I. (1985). From intention to actions: A theory of planned behavior. In J. Kuhl and J. Beckmann (Eds.), *Action-control: From Cognition to Behaviour*. Heidelberg: Springer. 11-39.
- Ajzen, I. (1988). *Attitudes, Personality and Behavior*. Milton Keynes: Open University Press.
- Ajzen, I. (1991). The Theory of Planned Behavior. *Organizational Behavior and Human Decision Processes* 50: 179-211.
- Batterman, R. (2000). Multiple Realizability and Universality. *British Journal for the Philosophy of Science* 51: 115-145.
- Batterman, R. (2002). Asymptotics and the Role of Minimal Models. *British Journal for the Philosophy of Science* 53: 21-38.
- Batterman, R. (2011). Emergence, Singularities, and Symmetry Breaking. *Foundations of Physics* 41: 1031-1050.
- Bechtel, W. (2005). The Challenge of Characterizing Operations in the Mechanisms Underlying Behavior. *Journal of the Experimental Analysis of Behavior* 84:313-325.
- Bechtel, W. (2007). Reducing Psychology While Maintaining its Autonomy Via Mechanistic Explanations. In Schouten, M. and De Joong, H.L. (Eds.), *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction*. Blackwell Publishing. 172-198.
- Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Lawrence Erlbaum Associates.
- Bechtel, W. and Abrahamsen, A. (2007). Mental Mechanisms, Autonomous Systems, and Moral Agency. *Proceedings of the Cognitive Science Society* 95-100.
- Bogen, J. (2005). Regularities and Causality: Generalizations and Causal Explanations. *Studies in History and Philosophy of Science Part C: Studies in the History and Philosophy of Biological and Biomedical Science* 36: 397-420.
- Bunge, M. (2003). *Emergence and convergence: Qualitative novelty and the unity of knowledge*. Toronto: University of Toronto Press.
- Busemeyer, J.R., and Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review* 100 (3): 432-459.
- Chemero, A. and Silberstein, M. (2008). After the Philosophy of Mind: Replacing Scholasticism with Science. *Philosophy of Science* 75 (1): 1-27.
- Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy* 78: 67-90.
- Craver, C. (2006). When Mechanistic Models Explain. *Synthese* 153 (3): 355-376.
- Craver, C. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon Press.
- Darden, L. (2006). *Reasoning in biological discoveries*. Cambridge: Cambridge University Press.
- Dennett, D. (1971). Intentional Systems. *Journal of Philosophy* 68 (4): 87-106.

- Dennett, D. (1987). *The Intentional Stance*. Cambridge, Massachusetts: The MIT Press.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Mass: MIT Press.
- Eliasmith, C. (2010). How we ought to describe computation in the brain. *Studies in History and Philosophy of Science Part A*, 41: 313-320.
- Eliasmith, C., Stewart, T., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A Large-Scale Model of the Functioning Brain. *Science* 338, 1202.
- Field, H. (1978). Mental Representation. *Erkenntnis* 13: 9-61.
- Floridi, L. (2005). Is Semantic Information Meaningful Data? *Philosophy and Phenomenological Research* 70 (2): 351-370.
- Fodor, J. (1975). *The Language of Thought*. New York: Crowell Press.
- Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, Massachusetts: MIT Press.
- Fodor, J. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- Frigg, R. and Hartmann, S. (2009). Models in Science. In E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2009 Edition)*. Retrieved October 3, 2011, from <http://plato.stanford.edu/archives/sum2009/entries/models-science>.
- Glennan, S. (2005). Modeling Mechanisms. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (2): 375-388.
- Godfrey-Smith, P. (2000). Information, Arbitrariness, and Selection: Comments on Maynard Smith. *Philosophy of science* 67 (2): 202-207.
- Godfrey-Smith, P. (2004). On Folk Psychology and Mental Representation, in H. Clapin, P. Staines & P. Slezak (Eds.), *Representation in Mind: New Approaches to Mental Representation*. Amsterdam: Elsevier Publishers. 147-62.
- Godfrey-Smith, P. (2005). Folk Psychology as a Model. *Philosophers' Imprint*, 5 (6).
- Harris, S., Sheth, S., & Cohen, M. (2008). Functional Neuroimaging of Belief, Disbelief, and Uncertainty. *Annals of Neurology* 63 (2): 141-147.
- Huebner, B. (2011). Minimal Minds. In T. Beauchamp and R.D. Frey (Eds.), *Oxford Handbook on Animals*. Oxford University Press. 441-468.
- Ji, S. (1997). Isomorphism between cell and human languages; Molecular biology, bioinformatics and linguistic implications. *BioSystems* 44: 17-39.
- Ji, S. (1999). The linguistics of DNA: words, sentences, grammar, phonetics, and semantics. In *Molecular Strategies in Biological Evolution*, Volume 870 of the Annals of the New York Academy of Science: 411-417.
- Kaplan, D. and Bechtel, W. (2011). Dynamical Models: An Alternative or Compliment to Mechanistic Explanations? *Topics in Cognitive Science* 3: 438-444.
- Lettvin, J., Maturana, H., McCulloch, W., & Pitts, W. (1959) What the Frog's Eye Tells the Frog's Brain. *Proceedings of the Institute of Radio Engineers* 47: 1940-1951.
- Longino, H. (2006). Theoretical Pluralism and the Scientific Study of Behavior. In S. Kellert, H. Longino, and C.K. Waters (Eds.), *Scientific Pluralism*. Minneapolis: University of Minnesota Press. 102-132.
- Machamer, P., L. Darden, and C. Craver. (2000). Thinking About Mechanisms. *Philosophy of Science* 67 (1): 1-25.
- Maibom, H. (2003). The Mindreader & the Scientist. *Mind & Language*, 18: 296-315.
- Maibom, H. (2009). In Defence of (Model) Theory Theory. *Journal of Consciousness Studies*, 16 (6-8): 360-378(19).

- Marr, D. (1982). *Vision: A computation investigation into the human representational system and processing of visual Information*. San Francisco, CA: Freeman.
- Maynard Smith, J. (2000). The Concept of Information in Biology. *Philosophy of Science* 67 (2): 177-94.
- Parisien, C. and Thagard, P. (2008). Robosemantics: How Stanley the Volkswagen Represents the World. *Minds & Machines* 18: 169-178.
- Sarkar, S. (1996). Biological Information: a skeptical look at some central dogmas of molecular biology. In S. Sarkar (Ed.), *The Philosophy and History of Molecular Biology: New Perspectives*. Dordrecht: Kluwer. 187-231.
- Sarkar, S. (2000). Information in Genetics and Developmental Biology: Comments on Maynard Smith. *Philosophy of Science* 67 (2): 208-213.
- Sarkar, S. (2004). Genes Encode Information for Phenotypic Traits. In Christopher Hitchcock (Ed.), *Contemporary Debates in Philosophy of Science*. Malden, MA: Blackwell Publishing. 259-274.
- Searle, J. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences* 3: 417-457.
- Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, Mass: The MIT Press.
- Sehon, S. (1997). Natural-Kind Terms and the Status of Folk Psychology. *American Philosophical Quarterly* 34 (3): 333-344.
- Sereno, M. I. (1991) Four Analogies Between Biological and Cultural/Linguistic Evolution. *Journal of Theoretical Biology* 151: 467-507.
- Sterelny, K. (2000). “The “Genetic Program” Program: A Commentary on Maynard Smith on Information in Biology”. *Philosophy of Science* 67 (2): 195-201.
- Sterelny, K. & Griffiths, P. (1999). *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: University of Chicago Press.
- Sterelny, K., Smith, K. and Dickison, M. (1996). The Extended Replicator, *Biology and Philosophy* 11: 377-403.
- Stich, S. (1983). *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, Massachusetts: The MIT Press.
- Stocker, A. and Simoncelli, E. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience* 9, 578-585.
- Thagard, P. (2006). *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press
- Thagard, P. (2009). Why Cognitive Science Needs Philosophy and Vice Versa. *Topics in Cognitive Science* 1 (2): 237-254.
- Thagard, P. (2012). *The Cognitive Science of Science*. Cambridge, MA: The MIT Press.
- Thelen, E., and Smith, B. (1994) *A dynamic systems approach to the development of cognition and action*. Cambridge: MIT Press.
- Van Gelder, T., and Port, R. (1995). It’s about time: An overview of the dynamical approach to cognition. In R. Port and T. Van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press. 1-43.
- Weber, M. (2005). Genes, Causation and Intentionality. *History and Philosophy of the Life Sciences* 27: 407-420.
- Wilson, M. (2006). *Wandering Significance: An Essay on Conceptual Behaviour*. New York: Clarendon Press.
- Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings*. Cambridge, MA: Harvard University Press.