

## **Minds, models and mechanisms: a new perspective on intentional psychology**

Eric Hochstein

**Abstract:** In this article, I argue that intentional psychology (i.e. the interpretation of human behaviour in terms of intentional states and propositional attitudes) plays an essential role in the sciences of the mind. However, this role is not one of identifying scientifically respectable states of the world. Rather, I argue that intentional psychology acts as a type of phenomenological model, as opposed to a mechanistic one. I demonstrate that, like other phenomenological models in science, intentional psychology is a methodological tool with its own benefits and insights that complements our mechanistic understanding of systems. As a result, intentional psychology's distinctive scientific benefit is its ability to model systems in unique, non-mechanistic, ways. This allows us to generate predictions that we cannot otherwise generate using the mechanistic models of neuroscience and cognitive psychology necessary for various scientific tasks.

**Keywords:** intentional psychology; mechanistic model; phenomenological model; prediction; explanation

### 1. Introduction

Intentional psychology, sometimes called 'folk' psychology, is the interpretation and explanation of human behaviour in terms of the attribution of intentional states or propositional attitudes, such as *beliefs*, *desires* and *intentions* (as well as other mental representations). Some argue that as a scientific

explanation of human behaviour, intentional psychology ought to be replaced by more capable scientific theories from domains such as neuroscience and neurophysiology (Churchland 1981; Stich 1983; Ramsey, Stich, and Garon 1990, and more recently Bickle 2003). Others argue that it is essential to understand human cognition in virtue of identifying internal causal states of systems (most famously Fodor 1975, 1987, 2008, but see also Putnam 1967 and Field 1978). In this article, I demonstrate that intentional psychology *is* essential to our understanding of human cognition, but not in virtue of identifying scientifically respectable states of the world. Instead, I propose a characterization of intentional psychology as a phenomenological model, and demonstrate that it is an indispensable *complement* to the mechanistic explanations of behaviour provided by neuroscience and neurophysiology.

Phenomenological models are commonly used in science to generate predictions, as well as perform other tasks, that cannot be accomplished using mechanistic models which identify the structure and causal features of systems. By analogy, intentional descriptions work in conjunction with mechanistic descriptions in our scientific study of the mind, each with a different scientific burden to bear. Both mechanistic and phenomenological models have an integral role to play in scientific inquiry, and neither can be eliminated from scientific methodology in favour of the other.

In the next section, I provide an outline of the distinction between mechanistic and phenomenological models in science. In Section 3, I provide a brief sketch of the arguments traditionally raised against the scientific merits of intentional psychology, and show how interpreting it as a phenomenological model overcomes these arguments. Finally, in Section 4, I demonstrate that intentional psychology, understood in this way, is an essential tool in our scientific study of the mind. It is neither threatened by, nor an extension of, our mechanistic accounts of human behaviour. Instead, like other phenomenological models in science, it is a methodological tool with its own benefits and insights that complements our mechanistic interpretations of systems.

## 2. Mechanistic models versus phenomenological models

### 2.1. Mechanistic models

Models serve a variety of purposes within science, from hypothesis testing to analyzing data. One critical purpose of models, however, is to provide explanations. The sorts of models we use to provide explanations often depends on the sorts of questions that we are interested in answering, and the scientific domain in which we are working. When trying to provide explanations in scientific disciplines pertaining to the mind – disciplines such as neuroscience, cognitive science and cognitive psychology – the type of model typically considered ideal for explanations is a *mechanistic* model (Machamer, Darden, and Craver 2000; Bechtel 2005, 2008; Glennan 2005; Craver 2006; Wimsatt 2007; Thagard 2009, 2012; Eliasmith 2010). Mechanistic models provide explanations by demonstrating how objects within a system interact so as to produce some phenomenon. A rather straightforward example, offered by Machamer et al. (2000, p. 3), is the phenomenon of chemical neurotransmission:

in the [case] of chemical neurotransmission, a presynaptic neuron transmits a signal to a postsynaptic neuron by releasing neurotransmitter molecules that diffuse across the synaptic cleft, bind to receptors, and so depolarize the post-synaptic cell.

In order to provide a complete explanation of chemical neurotransmission, we must understand the various aspects of the mechanism that produces it. First and foremost, we must identify the entities involved in the production of the phenomenon. In this case, the entities would be the presynaptic neuron, the post-synaptic neuron, the neurotransmitter molecules, the receptors and the post-synaptic

cell. Similarly, we must identify how these entities are situated spatially and temporally within the mechanism. This includes identifying the 'relative locations, shapes, sizes, orientations, connections, and boundaries of the mechanisms components' (Craver and Bechtel 2006, p. 470). The way in which the parts are structured and placed within the mechanism determines how they can interact with one another in order to produce the relevant phenomenon.

Next, we must identify how the different parts work together. In the above example, the presynaptic neuron *transmits* a signal by *releasing* neurotransmitter molecules. These molecules *diffuse* across the synaptic cleft and *bind* to receptors, *depolarising* the post-synaptic cell. These interactions bring about the phenomenon we are trying to explain (the neurotransmission of chemicals).

A mechanistic model is, therefore, explanatory when it identifies four major components: The *phenomenon* (the thing in need of explaining), the *parts* (the relevant components within the system), the *organization* (the relevant way in which the parts are situated spatially and temporally within the system) and the *activities* (the interactions and processes that go on between the component parts in order to produce the phenomenon). By identifying these four components, mechanistic models can explain how the phenomenon is being produced by the system. With this rough sketch of mechanistic models in hand, we now have a general set of criteria with which to distinguish mechanistic models from non-mechanistic ones.

## 2.2. Phenomenological models

Providing mechanistic explanations is a common goal in the behavioural sciences, but it is not the only goal. We use non-mechanistic models for a variety of different purposes in the sciences of the mind. They can describe phenomena in need of explanation, measure or calculate crucial quantities, make predictions, summarize data and function as heuristics for designing experiments (Bogen 2005, p. 401;

Craver 2006, p. 355). Such models primarily characterize or describe some phenomenon without attempting to decompose it into parts and operations for better understanding. These *phenomenological* models play an important role in science different from their mechanistic cousins. Consider Hodgkin and Huxley's mathematical model of the action potential in the squid giant axon:

[This] model is derived in part from the laws of physics and chemistry, such as Ohm's law, Coulomb's law and the Nernst equation, and it can be used to derive myriad electrical features of many different kinds of neurons in many different species. Despite this accomplishment, the authors insist that their model is not an explanation. (Craver 2006, p. 356)

Despite the appearance of a mechanistic explanation, Hodgkin and Huxley's model makes no ontological commitments about the underlying mechanisms that produce the time course of the action potential in the squid giant axon. It merely characterises and predicts important electrical features of those neurons. Despite not being explanatory in the way many scientists studying the mind are interested in, the model has proven extremely important and predictive. Identifying mathematical dependencies is extremely beneficial, even if it does not mechanistically explain the phenomenon it describes. As another example, consider the use of a specific form of dynamic systems theory in cognitive science:

What came to be called dynamical systems theory (DST) enables investigators to visualize the change in the state of a system over time. The simplest case is a plot of the states traversed by a system through time, that is, the system's trajectory through state space. Each dimension of state space corresponds to one variable of the system, and each point in the space corresponds to one of the possible states of the system. (Bechtel 2008, p. 187)

By applying the appropriate set of differential equations, we can predict the trajectory of the system through this state space. Dynamical models have been used in cognitive science to describe and predict cognitive phenomena like the production of speech (Port 2003) and the movement of animals (Kelso 1995). What is important to note about such models, however, is that they often characterise and predict the behaviour of cognitive systems without appealing to underlying causal mechanisms. Instead, it is 'usually only observable behaviour [that] is mapped to the model' (Eliasmith 2010, p. 319). In this respect, these dynamical models act as phenomenological models, providing no mechanistic explanation for the behaviour of the system. Despite this, such models have played a significant role in our scientific study of cognitive systems.

It is also important to note that even though they do not provide *mechanistic* explanations, phenomenological models still provide some degree of explanatory power. Consider what it means to explain a mechanistic system: 'explanations 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). The use of phenomenological models still allows us to determine a limited range of counterfactual behaviours in virtue of being predictively adequate. In this respect, such models can provide us with limited explanations, even in contexts where they do not match the explanatory power provided by mechanistic models.

### 3. The role of intentional psychology in the behavioural sciences

#### 3.1. Reasons to doubt intentional psychology's scientific merits

A great deal of criticism has been raised over the past few decades regarding the scientific value of intentional psychology. Churchland (1981, pp. 68–76), for example, argues that intentional psychology is

an explanatorily deficient account of human action. Among the common mental phenomena that intentional psychology cannot explain are: mental illness, creative imagination, the nature and function of sleep, memory, perceptual illusions and the learning process (Churchland 1981, p. 73). In virtue of this, he proposes that it fails as an adequate theory of human action, and should be replaced.

Similarly, Stich (1983, p. 101) argues that even when intentional psychology works at explaining and predicting systems, it fails the moment we apply it to systems in exotic circumstances. The use of intentional psychology, for example, can explain and predict the behaviour of people in particular kinds of societies, but fails when we apply it to people in radically different types of societies. It therefore cannot be relied upon to be a consistent explanatory or predictive tool for scientific purposes.

A different sort of argument is that the attribution of intentional states to systems is an inherently *normative* practice. Put another way, when we describe a system using intentional psychology, we attribute to it a set of intentional states that we feel it *ought to have* if it were a rational agent in that scenario. This means that 'an element of interpretation, and dramatic interpretation at that, must be recognized in any use of the intentional vocabulary' (Dennett 1987, p. 342). As a result, attributions are not made based on any *empirical* facts about the structure of the system, since the same physical system can always be interpreted as having vastly different sets of intentional states (e.g. different sets of beliefs, desires, goals, etc.) that will be compatible with its behaviour. Thus a physical description of a system will always be compatible with numerous different intentional interpretations. Given this normative and interpretative feature of intentional descriptions, intentional psychology appears ill-suited to play a role in the empirical sciences that work explicitly with *physical* descriptions. Quine (1960, p. 218) seems to have this in mind when he tells us that 'the underlying methodology of the idioms of propositional attitudes contrasts strikingly with the spirit of objective science at its most representative'.

Finally, it is questioned whether we need intentional attributions in science at all. In the past, it was common for people to attribute intentional states to natural phenomena such as the wind, fire, rivers and mountains (Dennett 1969, p. 89; Churchland 1981, p. 74); yet, over time we have learned to easily and conveniently predict such phenomena far better and more accurately without such intentional terminology. A completed neuroscience may very well allow us to do the same when predicting human action (Churchland 1981; Stich 1983; Ramsey et al. 1990; Bickle 2003). So, how might intentional idioms be critical to scientific progress?

### 3.2. Intentional psychology as a phenomenological model

The first thing to note about intentional psychology is that it tells us very little about the underlying structure of systems. To speak of a system in terms of its *beliefs* or its *mental representations* does not tell us what mechanisms are generating its behaviour. As Dennett (1987, p. 48) notes, ‘we find that we all use [intentional] psychology knowing next to nothing about what actually happens inside people’s skulls’. But while intentional psychology may not identify *structural* features of systems, it can be remarkably good at *predicting* systems (Dennett 1971, 1987, 1991; Churchland 1981, pp. 72–73; Fodor 1987, p. 3; Clark 1988; Lahav 1992; Cummins 2000, p. 131; Bechtel 2007, p. 194). This means that our use of intentional psychology to predict systems does not require that we adopt any ontological commitments about the structure of the system we are predicting:

[D]eciding on the basis of available evidence that something is (to be treated as) an Intentional system permits predictions having a normative or logical basis rather than an empirical one, and hence the success of an Intentional prediction, based as it is on no particular picture of the

system's design, cannot be construed to confirm or disconfirm any particular pictures of the system's design. (Dennett 1971, p. 97)

Intentional psychology interprets and predicts the phenomenon of human action without providing a mechanistic explanation for it. This is *exactly* what phenomenological models do: 'They are complete black boxes: they reveal nothing about the underlying mechanisms and so merely 'save the phenomenon' to be explained' (Craver 2006, p. 360).

But if intentional psychology functions as a phenomenological model, then this would imply that intentional descriptions are not the sort of explanations that scientists studying the mind commonly seek, since they do not describe causal mechanisms. Yet, intentional states seem to be invoked in our explanations of human action all the time. I can explain Jane's absence from my party due to her being *angry* with me. Similarly, I might explain why my dog barked at a quick moving shadow in terms of his *believing* the shadow was an animal. In these cases, intentional descriptions play a large role in our explanations. Thus it appears, at least *prima facie*, that we do not treat intentional psychology as a type of phenomenological model.

I propose that this is not, in fact, a problem for the account I provide. Intentional psychology still provides us with *limited* explanations when searching for mechanisms, in exactly the same way phenomenological models provide us with a degree of explanatory power:

Because phenomenal models summarize the phenomenon to be explained, they typically allow one to answer some ['what-if-things-had-been-different']-questions. But [a more complete] explanation shows why the relations are as they are in the phenomenal model, and so reveals conditions under which those relations might change or fail to hold altogether. (Craver 2006, p. 358)

We see just this sort of thing when talking in terms of intentional states. To explain my dog's behaviour in terms of his beliefs does not show *why* his beliefs are as they are. Nor does it reveal *how* or *why* those beliefs change or cease to exist under different circumstances. But in virtue of being predictive, it allows us to explain in a limited fashion. In order to provide a more complete explanation of his behaviour, however, we must understand the internal neurological mechanisms that causally produce the regularities in behaviour that we *describe* in terms of intentional states. Intentional psychology identifies regularities in the system's interactions in the world (Dennett 1991; Bechtel 2007, p. 194), but does not describe the underlying structure of the system that produces them.

### 3.3. Responses to objections and criticisms

Given the limited explanatory power of phenomenological models in the life sciences, Churchland's objection to intentional psychology on the grounds that it lacks explanatory power loses some of its sting. The fact that intentional attributions cannot explain a number of mental phenomena is not grounds for elimination, since phenomenological models are often *not used to provide* such explanations.

Even though an intentional model posits mental states and objects that may not be neurologically plausible, this does not damage its usefulness as a predictive model. As Craver (2006, p. 361) tells us, 'for those who merely want to predict the target system's performance, biologically implausible simulations will work just fine'. It is only when we go in search of mechanistic explanations that such models become insufficient.

Viewing intentional psychology as a phenomenological model allows us to avoid other criticisms levelled against it as well. Consider Stich's criticism that intentional psychology fails as a predictive

model since it cannot be used to accurately predict a system in unusual situations. This limitation on applicability is not, in fact, a sign that intentional psychology fails as a useful scientific tool. On the contrary, this is a rather common characteristic of phenomenological models. Given that such models abstract away all the structural details of the system being modelled, they will often only be useful within a constrained set of situations:

One might, for example, build a model that is useful only within a narrow range of conditions (such as health, proper functioning or the absence of disturbing outside forces), but that fails outside of those narrow conditions. For example, one might provide a model of verb-tense generation that performs perfectly well when the brain and vocal cords are working properly, but that provides no insight into how the system will behave if something breaks or if the system is in extreme environmental conditions. (Craver 2006, p. 357)

As a phenomenological model, intentional attributions will work extremely well in particular sorts of environments and situations, just not others. This is not a sufficient reason for elimination, however, or else we would have to throw out all phenomenological models in science.

As for the objection that intentional psychology is inherently *normative* and thus inappropriate for use in empirical scientific practice, this assumes that the only types of descriptions that belong in our rigorous scientific study of the mind are those that explicitly identify the structure and causal interactions of systems. But this is simply untrue. Even if we grant that such descriptions are the sorts of descriptions we are searching for when explaining mental phenomena, this does not imply that *only* such descriptions are relevant to scientific inquiry involving the mind. If it did, then statistical models would be equally banished from our scientific methodology, since they too do not describe the physical constitution and causal interactions of systems. Instead, such models only 'focus on describing the

regularities in the data' that we collect about a given system, and so 'are silent with respect to the particular physical implementation' (Eliasmith 2010, p. 315). In this regard, they too involve an element of interpretation. Yet, such models have an undeniably crucial role in our scientific methodology. The normative aspect of intentional descriptions does not change the fact that intentional models are predictive. And it is *qua* predictive model that such descriptions are important to our scientific methodology, not *qua* empirical/normative model (as is the case with statistical models).

In fact, intentional characterisations bear a family resemblance to statistical characterisations. Both act as a type of phenomenological model. Phenomenological models come in all shapes and sizes after all. Certain types of dynamical models, as we saw, act as phenomenological models. Statistical models, meanwhile, act as different sorts of phenomenological models. Unlike dynamical models, statistical models predict by describing 'the probability of various measurable states of the system given other known states of the system' (Eliasmith 2010, p. 315). Similarly, intentional characterizations are simply another flavour of phenomenological model. They are relevantly similar to both statistical and dynamical models in virtue of remaining agnostic as to the physical implementation of the system being described and predicted. However, instead of basing predictions on dynamic systems theory, or probability theory, they predict based on the tenets of rationality.<sup>1</sup>

---

<sup>1</sup> Some might object to this suggestion of a family resemblance between intentional and statistical models. After all, while statistical models are based on the well-defined axioms of probability theory, the axioms of *rationality* (on which intentional attributions seem to be based) are not so apparent. It is questionable whether there even *are* any clearly definable axioms of rationality. When we predict the behaviour of others, we do so without any explicit understanding of what such axioms might be. In this respect, intentional models may be too different in kind from statistical models to warrant inclusion into the class of phenomenological models. The problem with this sort of objection is twofold. First, it is a mistake to think that all domains to which statistical models are applied are therefore axiomatised. That there are statistical axioms does not mean that any particular statistical model (e.g. of stock markets) is also axiomatised in the relevant respect. We do not, for instance, have the axioms of stock market behaviour just because we have a statistical model of the stock markets. You could, after all, have a statistical model of rationality. Second, and more importantly, this sort of objection is a red-herring. Even if we assume that there are no explicit axioms of rationality (which is still up for debate), it is hardly a necessary characteristic of phenomenological models that they be based on axioms. That was never the claim. Intentional models are, however, relevantly like statistical models in key respects: First, they are predictively valuable in scientific practice. Second, they make predictions without telling us structural or mechanistic details of the system. Third, they are often used to identify patterns and regularities in behaviour produced by mechanistic systems.

The fact that various types of phenomenological models predict in different ways means that each is useful in different contexts (Eliasmith 2010). The lack of explicit quantification makes intentional models ideal for predictions in contexts where quantification is either not possible, or difficult to determine. This makes them ideal for situations where statistical and dynamical models are unavailable and/or uninformative. Intentional models allow us to make predictions without having to quantify over features of the system that we may not know how to measure. In this regard, intentional models have their place among the ranks of other phenomenological models with their own role and benefits.

#### 4. Why settle for the phenomenological?

Even if we grant that intentional psychology can be characterised as a phenomenological model, the preceding discussion does seem to be dancing around the main issue: If a mechanistic model of human action is available to us through neuroscience or neurophysiology, do we really *need* intentional psychology for scientific predictions? In this section, I argue that the mechanistic models found in these domains are insufficient to fill the scientific role played by an intentional model.

##### 4.1. Phenomenological models can generate predictions that mechanistic models cannot

If we view models as scientific tools, then different tools may be necessary for different tasks and the model that is best for explanations will not necessarily be the best for predictions. And indeed, there are clear cases where phenomenological models outperform their mechanistic counterparts when it comes to predictions. Consider, for example, two different ways of modelling the behaviour of water. One

---

Finally, they are used in conjunction with mechanistic models to provide more complete understandings of systems. These similarities provide us with compelling reasons to consider intentional characterisations as a species of phenomenological model, just as statistical models are.

which models water based on its atomic constituents and how they causally interact to produce behaviour, and the other which says nothing about the mechanisms responsible for the behaviour and only describes behavioural regularities:

If one is studying diffusion or Brownian motion, one adopts a molecular perspective in which water is regarded as a collection of particles. [. . .] However, if one's concern is the behavior of water flowing through pipes, the best-fitting models are generated within a perspective that models water as a continuous fluid. Thus, one's theoretical perspective on the nature of water depends on the kind of problem one faces. Employing a plurality of perspectives has a solid pragmatic justification. There are different problems to be solved, and neither perspective by itself provides adequate resources for solving all the problems. (Giere 2006, p. 34)

In the case of water's movement through pipes, the use of a phenomenological model allows us to predict and describe the behaviour of water better than the mechanistic model. The model that is best for explaining why water moves (in terms of molecular motion) is not necessarily the ideal model for predictions. Similarly, there are compelling reasons to think that intentional models are better suited to certain scientific tasks than the mechanistic models we use in our neurological and physiological explanations.

To demonstrate, consider the stagnation in the neuroscience of vision that took place in the 1960s and 1970s. According to Marr, a strictly mechanistic understanding of vision proved unfruitful until neuroscientists realised the need to adopt a different perspective about the system, and what it was doing. More specifically, the neuroscience of vision was only able to proceed when neuroscientists began attributing intentional content to neurological mechanisms (Marr 1982; Bechtel 2008, pp. 24–28). Certain neurons, for example, were found to contain visual information *about* the edges and boundaries

of seen objects. This information, according to Marr (1982, p. 19), needed to be ‘analyzed and understood in a way that [was] independent of the particular mechanisms and structures that implement them in our head’. And as Dennett (1971, p. 90) rightly points out, the ‘possession of information or misinformation is just as Intentional a notion as that of belief’. Only by taking an intentional interpretation of the system, and attributing *intentional content* to neurological mechanisms could we proceed in making essential predictions. In this regard, intentional models go hand-in-hand with mechanistic models in our best scientific practice, each contributing to the study of the mind in ways the other is not well suited for.

#### 4.2. Phenomenological models are necessary for generating mechanistic models

Despite not explicitly describing mechanisms, intentional psychological descriptions are an essential tool in our scientific discovery and understanding of mechanisms. One of the primary uses of phenomenological models is precisely in helping to construct mechanistic ones. According to Bogen (2005, p. 401), phenomenological models are often used to ‘suggest and sharpen questions about causal mechanisms’, ‘suggest constraints on acceptable explanations’ and ‘support inductive inferences without which mechanisms could not be successfully studied’. In a similar vein, the attribution of intentional states to systems allows us to better understand ‘the task being performed by the [. . .] system, why it needs to be performed, and the constraints the task itself places on how it is performed’ (Bechtel 2008, p. 26). This provides us with insights into the physical implementation of the system being described.

With this in mind, Bechtel points out that there are two methods of intervening in a mechanistic system in order to learn about it. One involves going inside the system and altering the relevant parts and operations to see what results. The other involves altering only the environmental conditions in

which the system is placed, and observing the regularities in the overall behaviour caused by those environmental changes:

‘Each of these modes of intervention and detection provides different information about mental mechanisms. None of them alone provides a direct and complete account; rather, each is partial and potentially misleading’ (Bechtel 2008, pp. 38–39).

Describing the spatiotemporal parts and operations working inside a system very often tells us little about the possible environmental conditions in which the mechanism will be placed, and so cannot tell us the behavioural regularities of entire systems in different environmental situations:

The behaviour an entity exhibits is often dependent upon context and there is no reason to think that the account of an entity offered by any inquiry considers how it will behave under all conditions but only those which are the focus of inquiry. As engineers are well aware, how a component will behave when inserted into a particular kind of system often needs to be investigated empirically. (Bechtel 2008, p. 22)

Meanwhile, it is the very purpose of intentional models to characterise how entire systems behave in various sorts of situations and environments, and to identify behavioural regularities. According to Bechtel (2007, p. 10), intentional psychology ‘characterises regularities in the way cognitive agents respond to situations arising in their environment’. A similar point is made by Dennett (1987, p. 257), who claims that intentional psychology ‘is well-suited to describe, in predictive, fruitful and illuminating ways, the cognitive prowess of creatures in their environments’. The lesson to take away from this is that descriptions in terms of objects and operations must be supplemented by intentional descriptions

since both provide insights that the other lacks, and both play a key role in explaining and learning about mechanistic systems.

## 5. Conclusion

I have argued that intentional psychology is best understood as playing the role of a phenomenological model in our scientific study of the mind. Those who strive for a reductive or eliminative account of intentional concepts within science must be cautious. As a predictive tool, intentional attributions may play an indispensable role in helping to construct a complete mechanistic explanation of psychology, and thus are useful in ways that mechanistic models are not. A plurality of models are useful precisely because they do a plurality of things.

Similarly, even if we supposed that a completed mechanistic model could, in principle, be used in place of an intentional one, we rarely abandon useful and predictive models in science simply because we think a far more complex and unwieldy model gets the ontology more accurate. We talk about light in terms of 'rays' when studying optics, even though we do not think this accurately captures the ontology of light. We model electricity as though it were a continuous entity with 'flow' and 'current' when modelling electrical circuits, instead of a massive collection of electrons moving along a track (or as small clouds of probability as quantum mechanics suggests). We similarly use Newtonian physics in many cases instead of the more fundamental and accurate quantum mechanics. In a similar vein, intentional psychology (construed as a phenomenological model) has an important scientific role to play, and the construction of more detailed mechanistic models will not challenge that.

## Acknowledgements

A number of individuals were pivotal in the shaping of this study, and I would like to offer them my thanks. This includes: Chris Eliasmith, Micheal McEwan, Kurt Holukoff and Alexander Winthers. I am also very grateful for the feedback received on this article from the audience of the 2011 Epistemology of Modelling and Simulations Conference.

## References

- 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 *The Matter of the Mind: Philosophical Essays on Psychology*, eds. M. Schouten and H. Looren, Oxford: Basil Blackwell, pp. 172–198.
- Bechtel, W. (2008), *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, New York: Lawrence Erlbaum Associates.
- Bickle, J. (2003), *Philosophy and Neuroscience: A Ruthlessly Reductive Account*, Boston, MA: Kluwer Academic Publishers.
- 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.
- Churchland, P.M. (1981), 'Eliminative Materialism and the Propositional Attitudes', *The Journal of Philosophy*, 78, 67–90.
- Clark, A. (1988), 'Thoughts, Sentences and Cognitive Science', *Philosophical Psychology*, 1, 263–278.
- Craver, C.F. (2006), 'When Mechanistic Models Explain', *Synthese*, 153, 355–376.
- Craver, C., and Bechtel, W. (2006), 'Mechanism', in *Philosophy of Science: An Encyclopedia*, eds. S. Sarkar and J. Pfeifer, New York: Routledge, pp. 469–478.
- Cummins, R. (2000), 'How does it Work?' vs. 'What are the Laws?' Two Conceptions of Psychological Explanation', in *Explanation and Cognition*, eds. F. Keil and R. Wilson, Cambridge, MA: The MIT Press, pp. 117–145.
- Dennett, D.C. (1969), *Content and Consciousness*, New York: Humanities Press.
- Dennett, D.C. (1971), 'Intentional Systems', *Journal of Philosophy*, 68, 87–106.

- Dennett, D.C. (1987), *The Intentional Stance*, Cambridge, MA: The MIT Press.
- Dennett, D.C. (1991), 'Real Patterns', *The Journal of Philosophy*, 88, 27–51.
- Eliasmith, C. (2010), 'How We Ought to Describe Computation in the Brain', *Studies in History and Philosophy of Science Part A*, 41, 313–320.
- Field, H. (1978), 'Mental Representation', *Erkenntnis*, 13, 9–61.
- Fodor, J.A. (1975), *The Language of Thought*, New York: Crowell Press.
- Fodor, J.A. (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: The MIT Press.
- Fodor, J.A. (2008), *LOT 2: The Language of Thought Revisited*, Oxford: Clarendon Press.
- Giere, R. (2006), 'Perspectival Pluralism', in *Scientific Pluralism*, eds. S. Kellert, H. Longino, and C.K. Waters, Minneapolis: University of Minnesota Press, pp. 26–41.
- Glennan, S. (2005), 'Modeling Mechanisms', *Philosophy of Science Part C: Studies in the History and Philosophy of Biomedical Sciences*, 36, 375–388.
- Kelso, J.A.S. (1995), *Dynamic Patterns: The Self-Organization of Brain and Behavior*, Cambridge, MA: The MIT Press.
- Lahav, R. (1992), 'The Amazing Predictive Power of Folk Psychology', *Australasian Journal of Philosophy*, 70, 99–105.
- Machamer, P.K., Darden, L., and Craver, C.F. (2000), 'Thinking About Mechanisms', *Philosophy of Science*, 67, 1–25.
- Marr, D. (1982), *Vision: A Computation Investigation Into the Human Representational System and Processing of Visual Information*, San Francisco, CA: W.H. Freeman.
- Port, R. (2003), 'Meter and Speech', *Journal of Phonetics*, 31, 599–611.
- Putnam, H. (1967), 'Psychological Predicates', in *Art, Mind and Religion*, eds. W.H. Capitan and D.D. Merrill, Pittsburgh, PA: University of Pittsburgh Press, pp. 37–48.
- Quine, W.V.O. (1960), *Word and Object*, Cambridge, MA: The MIT Press.
- Ramsey, W., Stich, S., and Garon, J. (1990), 'Connectionism, Eliminativism, and the Future of Folk Psychology', *Philosophical Perspectives*, 4, 499–533.
- Stich, S. (1983), *From Folk Psychology to Cognitive Science: The Case Against Belief*, Cambridge, MA: The MIT Press.

Thagard, P. (2009), 'Why Cognitive Science Needs Philosophy and Vice Versa', *Topics in Cognitive Science*, 1, 237–254.

Thagard, P. (2012), *The Cognitive Science of Science*, Cambridge, MA: The MIT Press.

Wimsatt, W.C. (2007), *Re-Engineering Philosophy for Limited Beings*, Cambridge, MA: Harvard University Press.