PENULTIMATE DRAFT ONLY (21/1/2011).
COMMENTS WELCOME. PUBLISHED VERSION AVAILABLE UPON REQUEST.

Models as interpreters
(with a case study from pain science)

Chuanfei Chin

University of Oxford and National University of Singapore

phiccf@nus.edu.sg

Most philosophical accounts of scientific models assume that models represent some aspect, or some theory, of reality. They also assume that interpretation plays only a supporting role. This paper challenges both assumptions. It proposes that models can be used in science to interpret reality. (a) I distinguish these interpretative models from representational ones. They find new meanings in a target system's behaviour, rather than fit its parts together. They are built through idealisation, abstraction and recontextualisation. (b) To show how interpretative models work, I offer a case study on the scientific controversy over foetal pain. It highlights how pain scientists use conflicting models to interpret the human foetus and its behaviour, and thereby to support opposing claims about whether the foetus can feel pain. (c) I raise a sceptical worry and a methodological challenge for interpretative models. To address the latter, I use my case study to compare how interpretative and representational models ought to be evaluated.

scientific model; interpretation; representation; meaning; context; foetal pain

Acknowledgements

This paper was written while I was a doctoral scholar funded by the British Society for the Philosophy of Science. Earlier versions were presented in 2008 at the 'Model-based representation in scientific practice' workshop held in the National University of Singapore and at the Applied Ethics group in the University of Oxford's Uehiro Centre for Practical Ethics. I am grateful for comments from both audiences, especially those from Dominic Wilkinson and Julian Savulescu. An anonymous referee suggested that I make some important clarifications. The case study on foetal pain is partly drawn from my B.Phil. thesis, supervised by Katherine J. Morris and William Child. I thank the editor, Axel Gelfert, for his advice and invitation to contribute to this volume.

Models as interpreters

(with a case study from pain science)

1 Introduction

To survey the local scene: in philosophy of science today, two accounts of scientific models prevail. Both are sometimes seen as rivals, but it is possible to imagine a truce between them. In my version of the truce, the semantic account describes structures. It assumes that, when a model is used to test a theory, they must share some logical structure. We can describe this shared structure with the mathematics used in the theory or with a set-theoretical predicate defined by it. The autonomous account, on the other hand, describes processes. It looks more closely at how a model is built and used. It reminds us that a model may have other functions besides testing a theory; a model can even be used to change a theory's structure or construct a new theory. We can discern some logic too in these processes, though it is not obviously formalisable. This truce has the virtue of not slurring any account as shallow, scholastic or merely sociological. Other truces have been proposed.[1] They differ in how intellectual labour is divided between semantic and autonomous accounts, and how stable that division is meant to be.

I will not disturb these truces. Instead, I want to draw our attention to what is common to both accounts of scientific models. They assume that models *represent* reality. More precisely: models represent some aspect, or some theory, of reality. Hughes (1997) is explicit on this. 'The characteristic – perhaps the only characteristic – that all theoretical models have in common is that they provide representations of parts of the world, or of the world as we describe it' (p. 325). While Morrison and Morgan (1999) insist that models are autonomous agents with a variety of functions, they agree that a model's 'representative

---

[1] I have learnt from da Costa & French (2000) and Portides (2008). Hartmann (1995), which distinguishes static and diachronic views, now seems prescient.

power' is what enables us to investigate reality with it and learn from it. Giere (2004) emphasises the pragmatic aspects of modelling, but these relate to what he calls 'the basic representational relationship': 'Scientists use models to represent aspects of the world for specific purposes' (p. 742). In his survey of semantic and autonomous accounts, and others which came before them, Portides (2008) sums up this consensus. 'Philosophers have identified several kinds of models used in science…all of which are different means of representing respective target systems in idealized and abstract ways' (p. 385). A related assumption: interpretation, if mentioned at all, plays a supporting role. It is needed to prepare experimental evidence for modelling (Morrison & Morgan, 1999, p. 31) or to translate a model's theoretical results into real terms (Hughes, 1997, p. 32).

This paper challenges both assumptions. I propose that models can be used in science to *interpret* reality, not just to represent it. In such models, interpretation plays the leading role.[2] I will distinguish these interpretative models from representational ones. They are designed to find new meanings in a target system's behaviour, not to fit its parts together. To explain how their aims and tools differ, I will use an example from political science. Then I will offer a case study from pain science. This case study describes the current scientific controversy on foetal pain. It highlights one source of this controversy: pain scientists with opposing perspectives are using different models to interpret the human foetus and its behaviour. I will use this case study to explore two further aspects of interpretative modelling: Does it face any special risk? How can scientists and others evaluate its success? My aim is neither to draw a sharp line between representation and interpretation in scientific practice, nor to deny that some models can be used for both. It is rather to highlight those aspects of interpretation that are obscured by an emphasis on representation.

---

[2] Two precursors: Nersessian (1999) describes models as 'interpretations of target physical systems, processes, phenomena, or situations'. Bailer-Jones (2002) offers a metaphorical approach that treats a model as 'an interpretation of an empirical phenomenon' (p. 124). More distantly, some logical positivists saw models as interpretations of a theory's abstract calculus – with aesthetic, didactic or heuristic value; see da Costa & French (2000), p. 117.

2 Representational and interpretative models

Let me distinguish two kinds of models, representational and interpretative. I want to clarify their different aims, tools and criteria for success. The representational model is more familiar. Scientists use it to represent a target system. The aim: to exploit an internal dynamic in the model in order to explain some data about the target system's behaviour.[3] This form of explanation relies on an analogy between the model's internal dynamic and the way in which parts of the target system fit together and interact. An internal dynamic may describe some causes or mechanisms of the target system (Glennan, 2005). Or it may, less ambitiously, be a set of equations that describes some structural dependencies in the target system (Morrison, 1999). How well a model represents is partly dependent on how much it can explain by selecting and then fitting parts together. Take the billiard ball model of a gas in Figure 1. Like many models, it mediates between theory and data. It uses a set of moving balls to represent the gas. The random movements of these balls form an internal dynamic. While this dynamic does not describe the movements of individual molecules within the gas, it can account for some data about how the gas behaves.

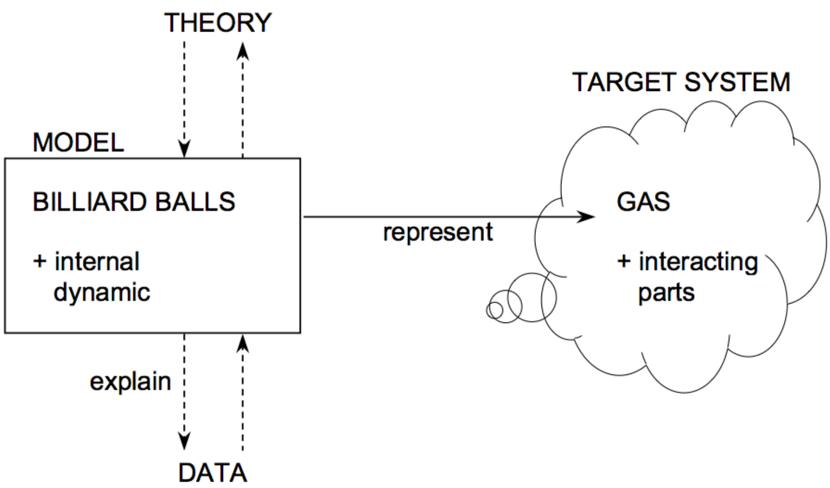

Fig. 1 *Analysis of a billiard ball model*

[3] The term 'internal dynamic' is introduced in Hughes (1997).

Two tools are often used to build a representational model: idealisation and abstraction (Frigg & Hartmann, 2006; Portides, 2008). Idealisation in a model distorts some properties of the target system, while abstraction leaves some properties out of the model. Both tools simplify the inner dynamic in the billiard ball model. First, the balls are assumed to move in perfectly straight lines at constant speeds until they hit the walls of a container. Second, ball sizes are omitted from the model.

A representational model's success is judged by many criteria. I will list only some common ones. First, scientists look at a model's empirical adequacy. This assesses how much of the original data about the target system's behaviour is explained by a model's internal dynamic. Second, they ask if the model makes successful predictions after manipulation. Both new and novel predictions may be involved (Hughes, 1997, p. 332). New predictions stay within what scientists previously expected about the target system's behaviour, while novel predictions stray into behaviour that was previously unsuspected. Third, they consider the model's simplicity and tractability. The latter is partly gauged by how well a model is adapted to the scientists' experimental and mathematical techniques. Fourth, they see if the model can be successfully integrated with nearby models or within general theories. These criteria may be juggled against each other when creating a model; and they may be weighed together when comparing models. I do not know of any hard and fast rules for doing so.

How is an interpretative model different? It is used to interpret a target system. Such a model does not explain data about a target system's behaviour by fitting its parts together. Rather, it makes sense of the data by finding new meanings in the system's behaviour. I find this contrast easier to explain with an example from political science. Consider, in Figure 2, the patriarchal model of kingship, which sees kings as the 'fathers of their people'.[4] This model, developed into its modern form in the sixteenth and seventeenth centuries by thinkers such as Jean

[4] The phrase is from Robert Filmer's *Patriarchia*, discussed in Schochet (1975) and Seidler (1993). I am vastly simplifying a sophisticated theory of political authority – not least by calling it an example from a science.

Bodin and Robert Filmer, tries to make sense of two strange aspects in a king's behaviour. (a) He rules without explicit consent from the people. (b) He has some control over their moral and religious lives. These data take on new meanings when interpreted by the patriarchal model. (a) People do not choose their king just as a son does not choose his father. Their subservient relationship is thus divinely ordained and rightful. (b) A king has some control over morality and religion within his realm because he has a duty to protect and improve his people morally; this is analogous to a father's duty to be a moral guide. Unlike a representational model, no analogy is drawn here in terms of internal dynamics. The analogy operates instead at the level of two new meanings: natural hierarchy and duty.

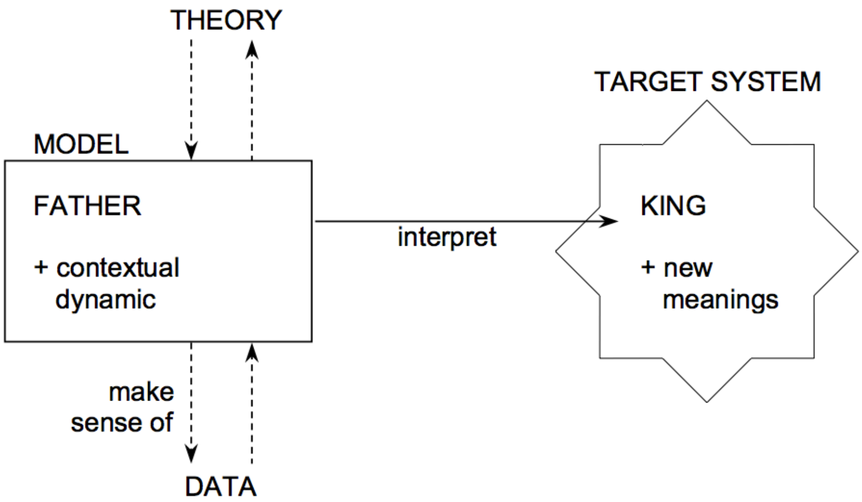

Fig. 2 *Analysis of the patriarchal model*

Some idealisation and abstraction may be useful in interpretation. For instance, a king's dubious character may have to be idealised when he is modelled as a fatherly guide. In Filmer's patriarchal model, part of the king's authority is traced back to Adam's status as the 'first father'. This interpretation holds only if the model abstracts away from the messy details of interrupted succession from Adam onwards (Seidler, 1993, p. 562). I think a third tool, recontextualisation, is even more important. When a model is used to interpret a target system, it brings that system into a new context, which comes with its own set of meanings. Data about the target system's behaviour are then interpreted in this new context, so that some of the meanings associated with the new context can be attributed to the system's behaviour. In the patriarchal model, the king is put into a new context: the natural hierarchy ordained by God. Within this context, the king's unelected power can be seen as analogous to other natural and unchosen forms of authority. To coin a convenient phrase: we might say that an interpretative model exploits a *contextual dynamic*, rather than an internal one.

Thus far, I have contrasted the aims and tools of representational and interpretative models. I have also listed some criteria of success for representational models. Next, I apply my analysis of interpretative models to a case study in pain science. This will show that interpretative models are used in science. Since the scientific controversy in this case study is not widely known, I will describe its background and arguments. This will prove valuable in two other ways: it uncovers a special risk faced by interpretative models, and helps to identify some criteria of success for them.

3 The foetal pain controversy

My case study is about the scientific controversy over foetal pain.[5] I begin by defining the scientific problem of foetal pain and clarifying its relation to three problems raised in philosophy about other minds. Then I discuss the conceptual and empirical arguments about the human foetus that have led to an impasse between pain scientists. To explain part of this impasse, I suggest that these scientists are using conflicting models to interpret the human foetus and its behaviour. Finally, I reconstruct how one side in this impasse has developed a model of the foetus as an unborn patient.

---

[5] I will simplify some arguments in order to focus on a few key themes. For more background from a philosophical perspective, see Benatar & Benatar (2001). Policy perspectives are provided in *Critical care decisions in fetal and neonatal medicine: ethical issues* (2006), ch. 4, by the Nuffield Council on Bioethics, and *Scientific Developments Relating to the Abortion Act 1967* (2007), by the UK House of Commons Science and Technology Committee.

This is how the problem of foetal pain is usually described:

> Pain is a subjective experience. The foetus cannot tell us what it is feeling, and there is no objective method for the direct measurement of pain. To address the question of pain in the foetus, one must use indirect evidence from a variety of sources, and then make an informed guess (Glover & Fisk, 1999).

What is meant here by 'subjective experience'? The answer lies in how pain scientists define pain (Price, 1999; Glover & Fisk, 1999; Lee et al., 2005; Lowery et al., 2007). First, they agree that pain has an essential psychological component. It is a 'feeling' or sensation that requires some level of consciousness. Second, they distinguish pain from nociception. Nociception is the body's physiological reaction to noxious stimuli. It includes neural transmissions from stimulated receptors to the spinal cord, which lead to reflex reactions and further transmissions up to the brain. Nociception often leads to pain, but it is neither necessary nor sufficient for it. Some patients, for instance, feel pain without being disturbed by noxious stimuli; others withdraw reflexively from some noxious stimuli without feeling pain. Nociception clearly occurs in the human foetus. But when pain scientists raise the possibility of foetal pain, they are asking if it can experience the unpleasant feeling that often accompanies nociception.

Their problem is sometimes said to be a 'species of the other minds problem' in philosophy, albeit a 'particularly troubling one' 'in practice'. These 'arise in the same way' (Benatar & Benatar, 2001, p. 61). I think this needs clarification.[6] In philosophy, three different problems are posed about other minds. First, there is a scientific problem. This asks about the distribution of minds among creatures other than normal adult humans (Bayne, 2009). The problem of foetal pain can be seen as a version of it, as applied to human foetuses. Second, there is a sceptical problem. This takes for granted the concept of the mind, then asks:

---

[6] I thank an anonymous referee for pressing me to do this.

how can we justify the belief that other noisy and moving adult bodies have minds? Those who pose this problem usually share our common-sense belief that normal adult humans have minds. Their concern is to uncover a rational basis for this belief (Bayne, 2009; Hyslop, 2010). It is possible that a method to justify this common-sense belief may offer us a means to test a scientific hypothesis about the minds of foetuses. But even if we are wont to dismiss the sceptical problem as a philosopher's idle fancy, the problem of foetal pain remains urgent and unsolved.

Third, there is what is known in the literature as a conceptual problem (Avramides, 2000; Hyslop, 2010). This is a problem about our understanding of other minds, not about their distribution or justification. It asks: how can we understand the attribution of minds to others? As Nagel (1986, p. 19) puts it, 'How can I conceive of my own mind as merely one of many examples of mental phenomenal contained in the world?' This way of posing the problem does not take any concept of mind for granted. Scientists who pose the foetal pain problem do not go so far. They do not challenge our understanding of pain in normal adult humans. But they ask if the concept of pain we apply to normal adult humans has any relevance to human foetuses. In doing so, they cannot take the concept of pain for granted in at least this sense: while everyone agrees that pain requires some level of consciousness, this level has yet to be determined. It is, in fact, one source of conceptual confusion in their controversy (Glover & Fisk, 1999; Benatar & Benatar, 2001; Derbyshire, 2001).

3.1 The scientific impasse

In December 2006, a majority in the US House of Representatives voted in favour of the Unborn Child Pain Awareness Act. The vote was 250-162. This fell short of the two-thirds vote needed to pass it. It seeks to 'ensure that women seeking an abortion are fully informed regarding the pain experienced by their unborn child'. Doctors would be required to tell women that, at least 20 weeks after fertilisation, the foetus can feel pain. The act cites 'expert testimony' from pain science. It includes three 'findings' of fact to support its claim that, by 20 weeks after fertilisation, the foetus 'may experience substantial pain'.[7] First, that foetus is said to have the 'physical structures necessary to experience pain'. Second, the foetus draws away from noxious stimuli in a way 'which in an infant or an adult would be interpreted as a response to pain'. Third, anaesthesia is routinely used in surgery for foetuses of that age.

Earlier, in August 2005, the *Journal of the American Medical Association* published a review that tries to rebut these findings. Lee et al. (2005) deny that a foetus can feel pain before the third trimester. This controversy over foetal pain is not new; another debate, with structurally similar arguments, appeared nearly a decade before in the *British Medical Journal* (Derbyshire & Furedi, 1996; Glover & Fisk, 1996; Lloyd-Thomas & Fitzgerald, 1996). In both debates, neither the scientists who are more sympathetic to the idea of foetal pain – including its possibility before the third trimester – nor those who are more sceptical appear to convince each other. Although impasses are not unusual in science, I find the logic of this one fascinating. Both the sceptical and sympathetic perspectives constantly invoke the same data about foetal neurology, behaviour and physiology. But they interpret them to reach opposite conclusions.

Let me show how this impasse arises. Consider, first, their dispute over foetal neurology. From the sceptical perspective, Lee et al. (2005) argue that the foetus develops the necessary neural structures to experience pain only at 29-30 weeks after gestation. This argument rests on three claims: Pain perception requires 'conscious recognition or awareness' of noxious stimuli (p. 947). Conscious awareness of those stimuli, in turn, requires functional thalamocortical connections. Yet tests of pre-term neonates show constant electrical activity in these connections only at 29 weeks, followed by synchronous activity at 30 weeks.

[7] Details are in the *Unborn Child Pain Awareness Act* (2006). The experts are not named, but pro-life advocates in the US (including the act's legislative sponsors) often cite Wright (1996), Glover & Fisk (1999) and Anand (2004). For a survey of related legislation, see 'The science, law, and politics of fetal pain legislation' (2002) in the *Harvard Law Review*.

But why is this cortical processing necessary? Lee et al. (2005) simply assert, as many do, that activity in its cortex is necessary for the foetus to feel pain. They leap from the uncontested claim that pain perception in a foetus requires ‘conscious recognition’ of a noxious stimulus to the assumption that it requires ‘cortical recognition’ (p. 949).[8] Some sceptics about foetal pain, such as Derbyshire (2006), claim support from brain experiments which allegedly show that cortical activity is necessary for an adult to feel pain. These inferences from neuroscientific data are sometimes challenged on conceptual and empirical grounds (Chin, 2007; Lowery et al., 2007). But, more directly, their critics from the sympathetic perspective reject any assumption that foetuses must use the same neural structures as adults to perceive pain (Anand, 2006; Lowery et al., 2007).

Even if cortical processing is necessary, there are other pathways to transmit pain signals to the cortex. Lee et al. (2005) focuses on direct thalamocortical connections that begin to form at 23-24 weeks. But indirect connections between the thalamus and cortex form earlier through the cortical subplate. Monoamine fibres that bypass the thalamus completely reach the cortex at about 16 weeks (Glover & Fisk, 1999). Here is how Lee et al. (2005) dismiss subplate connections:

> no human study has shown that synapses between subplate and cortical plate neurons convey information about pain perception from the thalamus to the developing cortex. (p. 950)

This suggests an encouraging empiricism. Yet it is difficult to imagine how any experiment could provide the evidence they seek. Adults cannot be tested since they lack subplates. Scientists may be able to detect activity in the foetuses’ subplate connections. But it is not obvious how they can ‘show’ that these

---

[8] Other scientists share this assumption: Lloyd-Thomas & Fitzgerald (1996), Glover & Fisk (1996), Smith et al. (2000) and Mellor et al. (2005). McCullagh (1997) offers some rare dissent: ‘Two questions – whether the cortex is normally involved in the appreciation of pain and whether it is necessary for this – are regularly conflated’ (p. 302).

connections are conveying ‘information about pain perception’ without assuming at the same time that foetuses *can* feel pain. To anyone from the sceptical perspective, that will seem to beg the question.

Connections via the subplate are sometimes rejected by sceptics because it is a ‘transient’ foetal structure (Lee et al., 2005; Derbyshire, 2006). But those from the sympathetic perspective deny that this transience prevents pain perception. They insist that neurons in the subplate are already functional, forming an ‘intrinsic synaptic network’ (Anand, 2006; Lowery et al., 2007). Moreover, there is some evidence that those subplate neurons which remain in the adult’s deep cortex play vital roles in cognitive processing.

This stalemate over neurology leads both sides to their opposing interpretations of foetal behaviour and physiology. Their dispute is over three striking series of experimental data about foetuses and neonates of the same age. (a) They withdraw automatically from noxious stimuli. From 14 weeks onwards, the foetus moves away from touch (Glover & Fisk, 1996). By 26 weeks, pre-term neonates respond to noxious stimulation with measurable withdrawal reflexes (Andrews & Fitzgerald, 1994). (b) At 28-30 weeks, neonates respond to heel lancing with facial movements similar to those in adults suffering pain (Craig et al., 1993). These facial movements are absent during non-invasive procedures. (c) During invasive procedures, some foetuses show the same stress responses as infants and adults suffering pain. At 16 weeks, when a needle punctures a foetus’ hepatic vein through its innervated abdominal wall, its cerebral blood flow increases (Teixeira et al., 1999). This does not happen when the umbilical cord, which has no nerves, is punctured. Levels of the stress hormones, cortisol, β-endorphin and noradrenaline, also increase more when the needle punctures the hepatic vein. For the first two hormones, this change is seen in foetuses at 23 weeks (Giannakoulopoulos et al., 1994); for the third, this is seen as early as 18 weeks (Giannakoulopoulos et al., 1999).

From the sceptical perspective, Lee et al. (2005) argue that these behavioural and physiological responses cannot prove that the foetuses suffer pain.
(a) Withdrawal reflexes may only indicate nociception. When receptors at the skin detect a noxious stimulus, they transmit an impulse to motor neurons in the spinal cord. This triggers reflexive muscular contractions, moving the body away from the stimulus. No feeling of pain need be involved yet. Moreover, such reflexes are seen even in subjects who lack 'cortical function', including anencephalitic infants and persistently vegetative patients (p. 950). (b) Facial movements need not be 'cortically controlled' (p. 950). They are also made by neonates with significant cortical injury. (c) Stress responses during invasive procedures are mediated by the autonomic nervous system and the hypothalamic-pituitary axis without 'conscious cortical processing' (p. 951). Some hormonal changes also occur during non-painful activities such as exercise.
To those from the sympathetic perspective, such arguments will appear too dismissive. Lee et al. (2005) promise a 'systematic multidisciplinary review' of the foetal pain problem, with an 'evidence synthesis from various fields' (p. 947). But their sceptical arguments do not treat behaviour and physiology as independent lines of evidence for foetal pain. Complex bodily withdrawals, facial reactions, changes in cerebral blood flow and stress hormone levels – these are all put aside as reflexive or irrelevant for the same reason. They need not involve cortical processing. Lee et al. (2005) have imposed their requirement about neural structures on the rest of the data. Yet, as I have shown, this requirement is contestable from the sympathetic perspective.

Without this requirement, the same behavioural and physiological data can be interpreted differently. Let me take two examples of behavioural data. (a) Lee et al. (2005) dismisses the data about withdrawal reflexes. They cite studies showing that anencephalitic infants and persistently vegetative patients exhibit similar reflexes. But Anand (2006), a critic from the sympathetic perspective, notes that scientists who study these subjects disagree if they can feel pain. It is as much an open question as the problem of foetal pain. Lee et al. (2005) also cite Andrews and Fitzgerald (1994), who show that full-term neonates exhibit these reflexes at a threshold much lower than that for children or adults. This suggests that the capacity to distinguish noxious and non-noxious stimuli is still maturing at that age. But, to those from the sympathetic perspective, a maturing system need not be oblivious to pain. Some even interpret the lower thresholds in full-term and pre-term neonates to mean that they are more sensitive to pain (Craig et al., 1993; Hadjistavropoulos et al., 1997; Anand, 2001); perhaps they lack mature pain-inhibiting mechanisms.

(b) To dismiss the data about facial reactions, Lee et al. (2005) cite Oberlander et al. (2002) – a study showing that even preterm infants with parenchymal brain injury can display those reactions at 32 weeks. Lee et al. (2005) assume that such infants cannot feel pain due to severe cortical damage. Thus they conclude that facial reactions at any age cannot prove pain perception. But the original researchers draw the opposite conclusion. For Oberlander et al. (2002), these facial reactions indicate that 'even in the presence of a significant neurologic injury, preterm infants are able to mount a clinically recognizable pain reaction that does not differ from noninjured infants, and thereby can be assessed for the presence of pain' (p. 575).

According to my reconstruction, this impasse is quite general. Do complex bodily withdrawals, facial reactions and stress responses indicate that the foetus can feel pain? Those from the sceptical perspective, including Lee et al. (2005), Mellor et al. (2005) and Derbyshire (2006), deny so. All these behavioural and physiological responses are possible without cortical processing. Their inference: since cortical processing is necessary for pain perception, such responses cannot show that the foetus can feel pain. At best, they indicate nociception.

Sceptical perspective	Sympathetic perspective
cortical processing is necessary ↓ behavioural/physiological responses are nociceptive	subcortical processing may be sufficient ↑ behavioural/physiological responses may be expressive

Table 1 *Impasse over foetal behaviour and physiology*

Roughly speaking, one perspective's modus ponens is another's modus tollens. Their critics from the sympathetic perspective, including Anand (2004) and Lowery et al. (2007), believe that the same responses, taken together, may express the foetus' pain. They thereby infer that subcortical processing may be sufficient for pain perception. Both perspectives agree on the data about foetal neurology, behaviour and physiology; their impasse is over interpretation. It is no surprise then that this impasse is reflected in how each perspective interprets ordinary practices during foetal surgery. For sceptics, anaesthetics and analgesics inhibit foetal movement, control stress responses and improve surgical outcomes in the long term (Lee et al., 2005). Their critics agree – except that they believe some of these changes may come about because the drugs reduce foetal pain.

I have shown how the arguments about foetal pain, which are drawn from different scientific disciplines, lead to an interpretative impasse. To an outsider, one aspect of this impasse remains mysterious. Each perspective begins with a key presupposition, which it takes to be obvious and not in need of justification. Scientists from the sceptical perspective assume that cortical processing is necessary for the foetus to feel pain, while those from the sympathetic perspective believe that the foetus' behavioural and physiological responses may express its pain. What accounts for these different starting points? I will suggest one explanation using the idea of interpretative models.

3.1 Interpretative models of the foetus

My hypothesis: scientists from the sceptical and sympathetic perspectives are using different models to interpret the foetus. These interpretative models explain why they begin their arguments with different presuppositions about the foetus. I will sketch how two models – the Incomplete Adult and the Unborn Patient – may be at work in the scientific impasse over foetal pain. Then I will study the logical structure and development of the latter model within the sympathetic perspective.

Those from the sceptical perspective may be drawn to the Incomplete Adult model. This model adopts the traditional view that all stages in the human developmental process are successive approximations of the adult stage (Anand & Craig, 1996). The foetus is therefore interpreted as a 'little adult' with incomplete physical and psychological systems. It is an adult in this sense: to function in the same way as other adults, it needs the same neural structures that they use. As critics such as Anand (2006) and Lowery et al. (2007) observe, this assumption is crucial to many sceptical arguments. Using the Incomplete Adult model, scientists who believe that cortical processing is required for adults to feel pain can infer that it is also required for foetuses to feel pain. Some sceptics go further. Derbyshire (2006) argues that, since some cognitive and emotional inputs contribute to the adult experience of pain, they are necessary for foetuses to feel pain. Mellor et al. (2005) suggest that, for foetuses to feel pain, they need the same post-natal environment that adults have.

In contrast, scientists from the sympathetic perspective favour the Unborn Patient model. This model interprets the foetus as 'a patient in its own right' (Smith et al., 2000). This patient is unique because it is not yet born. Faced with biological and environmental limitations, it may be forced to express pain in its own way – through its behavioural and physiological responses to noxious stimuli. This is the key presupposition of the sympathetic perspective. Sceptics such as Mellor et al. (2005) point out that data and conclusions about pre-term neonates can be applied to foetuses of the same age only if we treat the foetus as 'an unborn newborn' (p. 456). From the sympathetic perspective, being unborn is not the same as being incomplete. Anand and Craig (1996) appeal to the idea of ontogenetic adaptation to suggest that the foetus is at a developmental stage which is complete 'in itself'. Its physical and psychological systems are not incomplete versions of the adult's. So the foetus need not share the same neural structures as the adult in order to feel pain (Anand, 2006; Lowery et al., 2007).

This Unborn Patient model is relatively new in western medicine. *Williams Obstetrics*, a key textbook, heralded it early in 1980: 'Who could have dreamed, even a few years ago, that we could serve the foetus as physician?'[9] The first international symposium on 'The Fetus as a Patient' was held in Yugoslavia in 1984. More than a decade later, in 1998, a clinical review of foetal medicine in the *British Medical Journal* noted a change in attitude: 'the most significant advance is that most professionals and parents consider the fetus as a separate individual and a potential patient in its own right' (James, 1998). It is striking how the scientific debate over foetal pain began only after this model of the foetus as a patient had taken root. The earliest scientific event I could find is a 1985 meeting in Britain on 'Pain and the Fetus' (Richards, 1985). It was a closed-door seminar for only twelve experts; much of the discussion concentrated on neonatal pain. A decade passed before the possibility of foetal pain was seriously debated in 1996 in the *British Medical Journal*.

I want to bring my earlier analysis of interpretative models to bear here. Consider the Unborn Patient model in Figure 3. This is not a representational model. Unlike a plastic Ziegler model, it does not exploit any internal dynamic in order to explain new experimental data about the foetus. Instead, it tries to make sense of these data by positing new meanings in the foetus' neurology, behaviour and physiology. The foetus' behavioural and physiological responses to noxious stimuli are interpreted as potential expressions of pain. Its subcortical processing is interpreted as being potentially sufficient for the perception of pain.

[9] This is quoted in Williams et al. (2001), which surveys how broader perceptions of the foetus have evolved. I will focus on developments in science and medicine. But, outside these fields, the idea of the foetus as an 'unborn child' has a long history. The Society for the Protection of Unborn Children was founded in Britain in 1966.

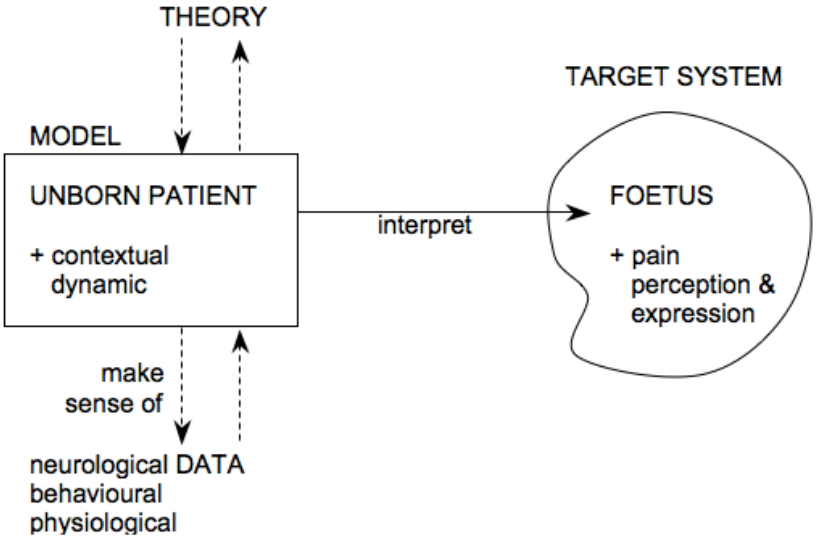

Fig. 3 *Analysis of the Unborn Patient model*

Again, recontextualisation plays an important role. The Unborn Patient model brings the foetus out of its old context, where it is seen as either an incomplete adult or an organ in a female adult (Wyatt, 2001). In its new context, the foetus is seen as a kind of patient. This model works by exploiting a contextual dynamic that includes other kinds of patient and the meanings associated with them. If it succeeds, then some of these new meanings, such as the capabilities to perceive and express pain, can be attributed to the foetus.

What is involved in this attribution? I will analyse two aspects of this process. First, there is conceptual expansion. Since the late 1980s, some scientists have developed a concept of neonatal pain which takes into account the neonate's limited ability to express its pain (Anand & Hickey, 1987; Hadjistavropoulos et al., 1997; Anand, 2001; Ranger et al., 2007). Others are extending this research to develop a similar concept of foetal pain. Their rationale: if a pre-term neonate can feel pain, so too might a foetus of the same age (Glover & Fisk, 1999; Smith et al., 2000; Anand, 2004). To support this conceptual expansion, they draw an analogy between the neonate's limited ability to express its pain and the foetus' predicament. 'The immature pain system thus uses the neural elements available during each stage of development to carry out its signalling role' (Anand, 2006).

Second, there is experimental discovery. Scientists have to collect data about how foetuses and neonates of the same age respond to noxious stimuli. Only then can their complex bodily withdrawals, facial reactions and stress responses be compared to those of infants and adults who are in pain. Analogy also operates here, partly at the level of experimental design. Scientists studying foetuses have mimicked crucial experiments done on neonates (Glover & Fisk, 1999). Anand et al. (1987) showed that, during surgery, neonates have hormonal stress responses similar to those of children and adults in pain. Such responses are reduced when the neonates are given fentanyl, a strong opioid analgesia.

These results convinced paediatric anaesthetists in Oxford to change their practice of not giving neonates analgesia or anaesthesia during invasive surgery (Richards, 1985); they also contributed to research developing a concept of neonatal pain. A team of scientists has now demonstrated that such results apply to foetuses too (Giannakoulopoulos et al., 1994; Giannakoulopoulos et al., 1999; Fisk et al., 2001). This mimicking recalls Kuhn's account (1996) of how research problems and techniques arise within normal science:

> What these have in common is not that they satisfy some explicit or even some fully discoverable set of rules and assumptions that gives the tradition its character and its hold upon the scientific mind. Instead, they may relate by resemblance and by modeling to one or another part of the scientific corpus which the community in question already recognizes as among its established achievements. (p. 45)

4 Evaluating interpretative models

I have described the aims and tools of interpretative models. Through my case study, I have also shown that these models are used, at least implicitly, in pain science. But I have said very little about their evaluation. This final section tackles that topic. I will distinguish two evaluative concerns about interpretative modelling: a sceptical worry and a methodological challenge. The sceptical worry can be defused, though it highlights a special risk faced by interpretative models. To meet the methodological challenge posed by this risk, I will return to my case study and propose some criteria for assessing the Incomplete Adult and Unborn Patient models.

The sceptical worry casts doubt on interpretative modelling in general. It asks: Can *any* interpretative model be imposed on the data?[10] I locate two sources of this worry. First, interpretative models are not constrained by data in the way that representational ones are. When scientists build a representational model, they are not free to choose any internal dynamic. They must choose one that can explain some data about the target system's behaviour. An interpretational model is not similarly constrained because its aim is not to explain those data. If this implies that any interpretative model can be imposed on the data, then interpretative modelling starts to look like an elaborate form of wishful thinking. Second, some may find interpretation to be generally suspicious. It seems to be neither rule-bound nor rigorous. It deals in meanings, which are famously slippery things. This is less troubling when interpretation plays only a supporting role in science, as in both semantic and autonomous accounts of scientific models. However, in my account of interpretative models, it plays a leading role.

I think the sceptical worry is unwarranted. Interpretative models are constrained by data too – but in a different way from representational ones. When scientists build an interpretative model, they cannot invent meanings as they wish. Their proposed meanings must make sense of some data about the target system's behaviour. That sets a standard of empirical adequacy which not every model can meet. To see what I mean, consider again the Unborn Patient model. As I emphasised earlier, its attribution of new meanings to the foetus depends on both conceptual expansion and experimental discovery. Through experiments,

[10] I am echoing Hesse (1966), who asks if '*any* scientific model can be imposed a priori on *any* explanandum and function fruitfully in its explanation' (p. 161). 'Imposed' is already a hostile term here. Similarly, a sceptic about interpretative models may challenge the claim that new meanings are 'found' in a target system's behaviour. I will sometimes speak, more neutrally, of meanings that are 'attributed'.

scientists from the sympathetic perspective looked for substantial similarities between the behavioural and physiological responses of the foetus, neonate and adult. But nature need not have cooperated. If it had turned out that foetuses do not have the same hormonal stress responses as those of neonates and adults in pain, or that such responses are not reduced when foetuses are given analgesia, then the Unborn Patient model would not yet be viable.

It is not obvious that representation, when used in models, is any more rule-bound or rigorous than interpretation. Nonetheless, I find the general suspicion about interpretation and meanings to be useful. It draws our attention to a special risk of interpretative models. Using an interpretative model, scientists posit an analogy between the meanings within the model's context and those in a target system's behaviour. The risk: a hasty attribution of meanings. A representational model does not face this risk. Scientists posit, instead, an analogy between the model's internal dynamic and the way in which parts of the target system fit together and interact. Their risk is a hasty generalisation of internal dynamics.

So how can the hasty attribution of meanings be avoided? Put another way: how can a model that aims to attribute meanings be assessed? That is the methodological challenge.[11] I take it seriously because scientists do not seem to have any explicit criteria for assessing their interpretative models. In the foetal pain controversy, they often accuse others of relying on models that lead to misleading assumptions. Mellor et al. (2005) criticises the tendency to interpret the foetus as an 'unborn newborn', which leads others to ignore the 'unique' conditions that the foetus lives in. Anand (2006) and Lowery et al. (2007), on the other hand, urges others not to treat the foetus as a 'little adult'; that model leads them to forget that structures and mechanisms for pain processing in

---

[11] As far as possible, I wish to stay neutral about the metaphysics of interpretation. My methodological challenge arises in different terms for those who believe that, through interpretation, new meanings are created and those who believe that they are discovered. Creators have to decide if their new meanings have been properly made, while discoverers have to determine if theirs have been correctly found. But it is not clear to me that these challenges differ in practice. I thank the anonymous referee for urging me to clarify this point.

foetuses are ‘unique’ and ‘different’ from those used by adults. Significantly, none of these scientists acknowledge that, to describe the foetus as ‘unique’ in their respective senses, they must be using other models to interpret it. Not surprisingly then, they do not offer any criteria for judging their own models. Using my case study of the Incomplete Adult and Unborn Patient models, I shall draw up a list of possible criteria for evaluating an interpretative model and compare them to the criteria for evaluating a representational model. These are rough heuristics, not strict rules. My list is exploratory; I hope to instigate more research by pointing out some difficulties with these criteria.

(a) Let me start with a criterion that has already been brought up: *empirical adequacy*. For a representational model, this assesses how much data about the target system is explained by the internal dynamic it proposes. For an interpretative model, this assesses something else: how much data is made intelligible by the new meanings it attributes. As I noted above, this is a necessary criterion. It ensures that interpretative models cannot be imposed arbitrarily on target systems. However, it may not be sufficient. For instance, in the foetal pain controversy, this criterion cannot help to decide between the Incomplete Adult and Unborn Patient models. Both models make sense of exactly the same data about the foetus’ neurology, behaviour and physiology.

(b) Might it help instead to look at how well a model makes sense of the data? Scientists who use the Unborn Patient model draw an analogy between the pain experiences of foetuses, neonates and adults. This analogy alone suffices to connect three surprising sets of experimental data about their similar responses to noxious stimuli. That is a measure of the model’s *analogical structure*. This is sometimes noted, though it is also implicit in the way scientists from the sympathetic perspective mention the three sets of data collectively (Anand & Hickey, 1987; Glover & Fisk, 1999; Smith et al., 2000). This structure is easy to miss when sceptics, such as Lee et al. (2005), criticise each set of data individually. Complex bodily withdrawals, facial reactions and stress responses: taken as separate lines of evidence, these can be dismissed as indicating only nociception. Taken together, as a body of evidence that has what scientists like to call ‘non-trivial structure’, they may be harder to ignore.

On the other hand, sceptics can point to the analogical structure of their Incomplete Adult model. This model draws an analogy between the neural needs of foetuses, neonates and adults for pain perception. For Lee et al. (2005), the analogy produces a uniform requirement of cortical processing, which can be used to judge all three sets of experimental data. Mellor et al. (2005) extend the model outwards by drawing an analogy between the environmental needs of foetuses, neonates and adults for pain perception. This leads to a new concept of ‘endogenous anesthesia’: it makes sense of some disparate data about the womb’s warmth and the presence of adenosine, neurosteroids and prostaglandins in utero. They infer, from this ‘increasing body of evidence’, that the foetus’ environment prevents it from waking up and suffering pain (p. 461). I do not know how we can weigh the analogical structures of these models against each other. Analogical structure can track the development of each model, but it is not yet clear if scientists can use it to choose between the Unborn Patient and Incomplete Adult models.

(c) Some scientists in my case study do not invoke analogical structure, even implicitly. They prefer to challenge a model’s *analogical appropriateness*. But I do not find this criterion helpful in resolving the impasse over foetal pain. Take the example of Mellor et al. (2005), who warn against ‘inappropriate extrapolation from the postnatal state to the fetal condition’ (p. 465). They call for greater ‘scientific rigor’. Their complaint: evidence for foetal wakefulness has been ‘based on how certain fetal responses “resemble” newborn sleep-wake behaviors’, as opposed to a ‘true determination of fetal wakefulness per se’ (p. 456). But, to show that the foetus is sleeping by default, they have to assume a resemblance in the wakeful neural patterns for foetuses and neonates. It is not obvious why their particular extrapolation about neurology should be considered appropriate while other extrapolations about behaviour and physiology from the sympathetic perspective are not. I suggest that criteria such as structure and appropriateness are not nuanced enough for describing and evaluating how analogies are used in interpretative models. The criterion of *simplicity* will not be decisive here too: one perspective's inappropriate analogy is likely to be another's simplifying ontology. We need new criteria that capture the way in which scientists, using rival interpretative models, cannot dislodge one analogy without deploying another.[12]

(d) Thus far, I have assumed that any new data will fit into the interpretative model being evaluated. What happens if new data, either from accidents or experiments, turn out to be anomalous? This diminishes the analogous structure of the model. But another criterion can come into play: I call it *elaborative strength*. This measures how well the new meanings in a model can be adapted to make sense of anomalous data. For example, Mellor et al. (2005) points to some data that may trouble scientists who use the Unborn Patient model. A sufficiently threatening stimulus, such as hypoxia or hypercapnia, triggers different reactions from foetuses and neonates. It will arouse and wake up a neonate. On the other hand, it suppresses arousal in a foetus; its breathing and bodily movements lessen, while blood flow is redistributed from its peripheral organs to the heart and brain. These new data seem to clash with the model's analogy between foetal and neonatal experiences. But its defenders might argue: because of their unique environment, foetuses must respond differently when faced with a sufficiently threatening stimulus. They have limited resources and room in the womb, so withdrawal into a less active state offers them the most feasible path.

This elaboration of the Unborn Patient model, in order to make sense of anomalous data, is *ad hoc*. But at least it appeals to a principle that is widely accepted even by scientists from the opposing sceptical perspective: 'Unlike the situation postnatally, arousal costs the fetus critically limited energy and for little advantage' (Mellor et al., 2005, p. 460). This model's elaborative strength will be

---

[12] Bradie (1999) makes a similar point about scientific metaphors.

greater than if its defenders resort to an *ad hoc* manoeuvre that simply quarantined the anomalous data.[13]

(e) I also propose *predictive success* as a criterion. But prediction works differently for representational and interpretative models. A representational model exploits its internal dynamic to explain data about the target system's behaviour. Scientists can manipulate this internal dynamic to make predictions about the system's behaviour. An interpretative model does not have such an internal dynamic. So how can it be used to predict? It has a contextual dynamic that makes sense of data about the target system's behaviour. Scientists may be able to use that contextual dynamic to make predictions about the system, though these will be characterised in terms of its new meanings.

To test the Unborn Patient model in this way, scientists need to make predictions about foetuses in terms of meanings such as pain perception and expression. One possible prediction: if a subject suffers repeated and prolonged painful experiences as a foetus, its psychological health as a child or an adult will be damaged. Some scientists who believe that pre-term neonates can suffer pain have developed a research programme to identify the long-term neurological, behavioural and psychological impact of pain suffered by these neonates (Grunau et al., 1994; Anand, 1998; Grunau, 2002; Grunau et al., 2009); their research may be expanded to include foetuses. But any prediction must be sufficiently entangled with the new meanings of the Unborn Patient model, so that scientists using the Incomplete Adult model could not easily account for it. Only then can predictive success help to decide between these two models. Otherwise, we shall have to weigh the elaborative strength of the Incomplete Adult model against the predictive success of the Unborn Patient model.

(f) Like representational models, interpretative ones can be judged by their *integrative success*. This measures how well an interpretative model can be

[13] Giere (1990), p. 223, gives an example of *ad hoc* corrections, based on known principles, to a representational model in physics. Hartmann (1995), quoting Lakatos, reminds us that temporary quarantines can be rational in scientific modelling.

integrated with nearby models or within more general theories. For instance, scientists who use the Unborn Patient model to attribute pain perception and expression to foetuses can draw support from those who attribute other psychological properties to them (Anand 2006). Hepper (1996) proposes a concept of foetal memory. Dirix et al. (2009) have recently identified, in foetuses, complex patterns of habituation to repeated stimulation; they interpret these to be rudimentary forms of foetal learning and memory. Such research may partly answer critics, such as Derbyshire (2006), who believe that more elaborate psychological inputs are necessary for the foetus to feel pain. It is, as yet, an open question whether the Unborn Patient model and these new concepts can support a general theory of foetal psychology.

5 Conclusion

My proposal: models can interpret reality. Such models explore what a system means, rather than how it works. In this paper, I challenged the common assumptions that models in science represent reality, and that interpretation plays a supporting role in them. I started with an intuitive distinction between representational and interpretative models. A representational model explains data about the target system's behaviour by fitting its parts together, while an interpretative model makes sense of the data by finding new meanings in the target system's behaviour. Using a case study of the scientific controversy over foetal pain, I showed how interpretative models can be developed and used in science. I identified the risk of hastily attributing meanings to target systems and suggested some criteria for evaluating interpretative models. Table 2 summarises my preliminary findings about the similarities and differences between representational and interpretative models. My case study from pain science is of some intrinsic interest. It describes an interdisciplinary impasse in which scientists disagree over how to integrate multiple lines of evidence; and it shows how the idea of interpretative models can be used to analyse such an impasse.

	Representational models	Interpretative models
Aims	To explain data about target system To fit parts of target system together	To make sense of data about target system To find new meanings in target system
Core analogies	Between model's internal dynamic and target system's interacting parts	Between meanings in model's contextual dynamic and meanings in target system
Tools	Idealisation and abstraction	Idealisation and abstraction Recontextualisation
Risks	Hasty generalisation of internal dynamics	Hasty attribution of meanings
Evaluative criteria	Empirical adequacy ? Simplicity and tractability ? Predictive success (for new and novel facts) Integrative success	Empirical adequacy Analogical structure and appropriateness Simplicity Elaborative strength Predictive success (using new meanings) Integrative success

Table 2 *Summary on representational and interpretative models*

Let me conclude by addressing two questions. They appear to be terminological, but are related to the philosophical significance of my proposal. Should interpretative models be considered a kind of representational model? Representation may be a broad enough church to include them. As Hughes (1997) noted, 'the concept of representation is as slippery as that of a model' (p. 325). I have no quarrel with this, so long as it is acknowledged that interpretation can play a leading role in some scientific models and that these models may have different aims, tools, risks and criteria of success. It is my hope that these will form a new locus for philosophical analysis. (a) We need a better language to describe and evaluate how some models attribute new meanings to target systems. We need to investigate, more carefully, how recontextualisation works in science. (b) We need some ways to assess and compare analogical structure and elaborative strength. If these criteria are applicable to representational models too, are they assessed and compared in the same ways?

(c) We also need to find out how widespread and significant this sort of modelling is. I have offered one brief example from political science and one case study from pain science. Both occur in situations shaped by moral and political interests. Are interpretative models more likely to be used in such morally and politically fraught situations?[14] This is possible, though there are suggestive parallels between the models I discuss and what others variously describe as: the 'prepared descriptions' of phenomena, which are made with 'a good deal of practical wisdom' (Cartwright, 1983, p. 134), the 'conceptualising' of evidence and theory 'into compatible terms' (Morrison & Morgan, 1999, p. 31) and the 'reconceptualization' of phenomena (da Costa & French, 2000, p. 121).

This leads me to ask: Should interpretative models be called conceptions instead? We need not fuel what Hartmann (1995) sees as the 'almost inflationary use of the term "model" in science'. But this proposal runs three risks. (a) We may neglect these conceptions, especially since models are increasingly seen as central to scientific practice. (b) We may obscure parallels between representational and interpretative models. Some of these are listed in Table 2. Others are found in narratives about model-building. Portides (2008), for instance, emphasises that 'theoretical principles, semi-empirical results, and experimental findings are blended together' to build a representational model. I have shown that the same blending is used to build the Unborn Patient model. (c) We may also obscure parallels with models in the humanities and social sciences, where interpretation is sometimes given more weight. These may yet offer us new resources for analysing those models in science which serve as interpreters.

---

[14] I owe this question to the anonymous referee. It suggests that interpretative models can sometimes be used, as an analytical tool, to highlight the moral and political interests at stake in scientific controversies.

References

Anand, K.J., (1998). Clinical importance of pain and stress in preterm neonates. *Biology of the Neonate*, 73(1), 1-9.

Anand, K. J. (2001). Consensus statement for the prevention and management of pain in the newborn. *Archives of Pediatrics & Adolescent Medicine*, *155*(2), 173-180.

Anand, K. J., & Craig, K. D. (1996). New perspectives on the definition of pain. *Pain*, *67*(1), 3-6; discussion 209-211.

Anand, K. J., & Hickey, P. R. (1987). Pain and its effects in the human neonate and fetus. *The New England Journal of Medicine*, *317*(21), 1321-1329.

Anand, K. J., Sippell, W. G., & Aynsley-Green, A. (1987). Randomised trial of fentanyl anaesthesia in preterm babies undergoing surgery: effects on the stress response. *Lancet*, *1*(8524), 62-66.

Anand, K. J. S. (2004). *Expert report of Kanwaljeet S. Anand, MBBS, DPhil. Testimony before the Northern District of the US District Court in California*. http://www.nrlc.org. (Accessed 11 February 2010)

Anand, K. J. S. (2006). Fetal pain? *Pain: Clinical Updates*, *XIV*(2).

Andrews, K., & Fitzgerald, M. (1994). The cutaneous withdrawal reflex in human neonates: sensitization, receptive fields, and the effects of contralateral stimulation. *Pain*, *56*(1), 95-101.

Avramides, A. (2000). *Other minds*. London: Routledge.

Bayne, T. (2009). Other minds. In T. Bayne, A. Cleeremans, & P. Wilken (Eds.), *The Oxford companion to consciousness* (pp. 490-492). Oxford: Oxford University Press.

Benatar, D., & Benatar, M. (2001). A pain in the fetus: toward ending confusion about fetal pain. *Bioethics*, *15*, 57–76.

Bradie, M. (1999). Science and metaphor. *Biology and Philosophy*, *14*(2), 159-166.

Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Clarendon Press.

Chin, C., 2007. *Predicting and projecting pain: a philosophical study of pain science*. B.Phil. thesis. University of Oxford.

da Costa, N., & French, S. (2000). Models, theories, and structures: thirty years on. *Philosophy of Science*, *67*, S116-S127.

Craig, K. D., Whitfield, M. F., Grunau, R. V., Linton, J., & Hadjistavropoulos, H. D. (1993). Pain in the preterm neonate: behavioural and physiological indices. *Pain*, *52*(3), 287-299.

(2006). *Critical care decisions in fetal and neonatal medicine: ethical issues*, Nuffield Council on Bioethics.

Derbyshire, S. W., & Furedi, A. (1996). Do fetuses feel pain? 'Fetal pain' is a misnomer. *BMJ: British Medical Journal*, *313*(7060), 795.

Derbyshire, S. W. (2001). Fetal pain: an infantile debate. *Bioethics*, *15*(1), 77–84.

Derbyshire, S. W. G. (2006). Can fetuses feel pain? *BMJ (Clinical Research Ed.)*, *332*(7546), 909-

912.

Dirix, C. E. H., Nijhuis, J. G., Jongsma, H. W., & Hornstra, G. (2009). Aspects of fetal learning and memory. *Child Development*, *80*(4), 1251-1258.

Fisk, N. M., Gitau, R., Teixeira, J. M., Giannakoulopoulos, X., Cameron, A. D., & Glover, V. A. (2001). Effect of direct fetal opioid analgesia on fetal hormonal and hemodynamic stress response to intrauterine needling. *Anesthesiology*, *95*(4), 828-835.

Frigg, R., & Hartmann, S. (2006). Models in science. In *Stanford encyclopedia of philosophy (Winter 2009 edition)*. http://plato.stanford.edu/entries/models-science/. (Accessed 30 January 2010)

Giannakoulopoulos, X., Sepulveda, W., Kourtis, P., Glover, V., & Fisk, N. M. (1994). Fetal plasma cortisol and beta-endorphin response to intrauterine needling. *Lancet*, *344*(8915), 77-81.

Giannakoulopoulos, X., Teixeira, J., Fisk, N., & Glover, V. (1999). Human fetal and maternal noradrenaline responses to invasive procedures. *Pediatric Research*, *45*(4 Pt 1), 494-499.

Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, *71*(5), 742-752.

Giere, R.N., 1990. *Explaining science: a cognitive approach*, London: University of Chicago Press.

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), 443-464.

Glover, V., & Fisk, N. (1996). Do fetuses feel pain? We don't know; better to err on the safe side from mid-gestation. *BMJ (Clinical Research Ed.)*, *313*(7060), 796.

Glover, V., & Fisk, N. (1999). Fetal pain: implications for research and practice. *British Journal of Obstetrics and Gynaecology*, *106*(9), 881-886.

Grunau, R. V., Whitfield, M. F., Petrie, J. H., & Fryer, E. L. (1994). Early pain experience, child and family factors, as precursors of somatization: a prospective study of extremely premature and fullterm children. *Pain*, *56*(3), 353-359.

Grunau, R. (2002). Early pain in preterm infants. A model of long-term effects. *Clinics in Perinatology*, *29*(3), 373-394, vii-viii.

Grunau, R. E., Whitfield, M. F., Petrie-Thomas, J., Synnes, A. R., Cepeda, I. L., Keidar, A., Rogers, M., et al. (2009). Neonatal pain, parenting stress and interaction, in relation to cognitive and motor development at 8 and 18 months in preterm infants. *Pain*, *143*(1-2), 138-146.

Grush, R. (2002). Cognitive science. In P. Machamer & M. Silberstein (Eds.), *The Blackwell guide to the philosophy of science* (pp. 272-289). Blackwell Publishers Ltd.

Hadjistavropoulos, H. D., Craig, K. D., Grunau, R. E., & Whitfield, M. F. (1997). Judging pain in infants: behavioural, contextual, and developmental determinants. *Pain*, *73*(3), 319-324.

Hartmann, S. (1995). Models as a tool for theory construction: some strategies of preliminary physics. In W. E. Herfel (Ed.), *Theories and models in scientific processes*, Poznan studies in the philosophy of science and the humanities (pp. 49-67). Amsterdam: Rodopi.

Hepper, P. G. (1996). Fetal memory: does it exist? What does it do? *Acta Paediatrica (Oslo, Norway: 1992). Supplement*, *416*, 16-20.

Hesse, M. B. (1966). *Models and analogies in science*. University of Notre Dame Press.

Hughes, R. I. G. (1997). Models and representation. *Philosophy of Science*, *64*, S325-S336.

Hyslop, A. (2010). Other minds. In *Stanford encyclopedia of philosophy (Winter 2009 edition)*. http://plato.stanford.edu/archives/win2009/entries/other-minds/. (Accessed 10 May 2010)

James, D. (1998). Recent advances: fetal medicine. *BMJ: British Medical Journal*, *316*(7144), 1580-1583.

Kuhn, T. S. (1996). *The structure of scientific revolutions* (3rd ed.). University Of Chicago Press.

Lee, S. J., Ralston, H. J., Drey, E. A., Partridge, J. C., & Rosen, M. A. (2005). Fetal pain: a systematic multidisciplinary review of the evidence. *JAMA: The Journal of the American Medical Association*, *294*(8), 947.

Lloyd-Thomas, A. R., & Fitzgerald, M. (1996). Do fetuses feel pain? Reflex responses do not necessarily signify pain. *BMJ (Clinical Research Ed.)*, *313*(7060), 797-798.

Lowery, C. L., Hardman, M. P., Manning, N., Hall, R. W., & Anand, K. J. S. (2007). Neurodevelopmental changes of fetal pain. In *Seminars in Perinatology* (Vol. 31, pp. 275–282).

McCullagh, P. (1997). Do fetuses feel pain? Can fetal suffering be excluded beyond reasonable doubt? *BMJ (Clinical Research Ed.)*, *314*(7076), 302-303.

Mellor, D. J., Diesch, T. J., Gunn, A. J., & Bennet, L. (2005). The importance of 'awareness' for understanding fetal pain. *Brain Research Reviews*, *49*(3), 455–471.

Morrison, M. (1999). Models as autonomous agents. In *Models as mediators: perspectives on natural and social science* (Vol. 52, pp. 38-65). Cambridge: Cambridge University Press.

Morrison, M., & Morgan, M. S. (1999). Models as mediating instruments. In *Models as mediators: perspectives on natural and social science* (Vol. 52, pp. 10–37). Cambridge: Cambridge University Press.

Nagel, T. (1986). *The view from nowhere*. New York: Oxford University Press.

Nersessian, N. J. (1999). Model-based reasoning in conceptual change. *Model-based reasoning in scientific discovery*, 5–22.

Oberlander, T. F., Grunau, R. E., Fitzgerald, C., & Whitfield, M. F. (2002). Does parenchymal brain injury affect biobehavioral pain responses in very low birth weight infants at 32 weeks' postconceptional age? *Pediatrics*, *110*(3), 570-576.

Portides, D. (2008). Models. In S. Psillos & M. Curd (Eds.), *The Routledge companion to philosophy of science* (pp. 385-395). Routledge.

Price, D. D. (1999). *Psychological mechanisms of pain and analgesia*. Seattle: IASP Press.

Ranger, M., Johnston, C. C., & Anand, K. J. S. (2007). Current controversies regarding pain assessment in neonates. *Seminars in Perinatology*, *31*(5), 283-288.

Richards, T. (1985). Can a fetus feel pain? *British Medical Journal (Clinical research ed.)*, *291*(6504), 1220-1221.

Schochet, G. J. (1975). *Patriarchalism in political thought: the authoritarian family and political speculation and attitudes, especially in seventeenth-century England*. Basic Books.

(2007). *Scientific developments relating to the Abortion Act 1967: twelfth report of session 2006-2007*, UK House of Commons Science and Technology Committee.

Seidler, M. (1993). Religion, populism, and patriarchy: political authority from Luther to Pufendorf. *Ethics*, *103*(3), 551–569.

Smith, R. P., Gitau, R., Glover, V., & Fisk, N. M. (2000). Pain and stress in the human fetus. *European Journal of Obstetrics and Gynecology*, *92*(1), 161–165.

Teixeira, J. M., Glover, V., & Fisk, N. M. (1999). Acute cerebral redistribution in response to invasive procedures in the human fetus. *American Journal of Obstetrics and Gynecology*, *181*(4), 1018-1025.

(2002). The science, law, and politics of fetal pain legislation. *Harvard Law Review*, *115*(7), 2010-2033.

(2006). *Unborn Child Pain Awareness Act (HR6099)*. 109th US Congress, 2nd Session. http://thomas.loc.gov. (Accessed on 30 March 2007)

Williams, C., Alderson, P., & Farsides, B. (2001). Conflicting perceptions of the fetus: person, patient, 'nobody', commodity? *New Genetics and Society*, *20*(3), 225.

Wright, J. A. (1996). *Statement of Jean A. Wright, MD, MBA, Associate Professor of Pediatrics and Anesthesia. Testimony before the US House Judiciary Committee, Subcommittee on the Constitution*. http://www.nrlc.org. (Accessed on 12 February 2010)

Wyatt, J. (2001). Medical paternalism and the fetus. *Journal of Medical Ethics*, *27 Suppl 2*, ii15-20.