

Chapter 1

Introduction

At the outset we humans know nothing, or not very much.

Sometime later, if things go well, we do know some things.

How does the transition occur? How do humans come to know about objects, actions, and minds?

This question belongs to a family of questions about the origins of mind that philosophers have been asking for a while. In a beautiful myth, Plato suggests that the answer is recollection. Before we are born, in another world, we become acquainted with all the truths we will ever know. Then we are involved in an unfortunate traffic accident and fall to earth, forgetting everything. But as we grow we are sometimes able to recall parts of what we once knew. So it is by recollection that humans come to know about objects, actions and minds.¹

How else could this happen? Since Plato, philosophers and psychologists have offered other stories. Some hold that knowledge is in some sense present at birth, or else that the concepts which make knowledge possible are already present at birth. Others suggest that concepts and knowledge are acquired through sensory experience, through learning to act, through training in language or through social interaction. None of these bold, seductive ideas is supported by much evidence. They are too difficult to test, or perhaps not even precise enough to test. And when we look at particular domains of knowledge in detail—for instance, when we look at how humans come to know about minds—we will discover complexities that seem to be incompatible with any one of the stories. While it would be fun to tour nativism, empiricism and other big ideas about the developmental origins of human knowledge, we are unlikely to make much progress if the last couple of millennia are any guide to the future. Let's try a different approach.

Start with the details. Take one domain of knowledge—knowledge of

¹ This is approximately the story in Plato's *Phaedrus*. I might have made up the bit about the traffic accident.

objects, say. What has been discovered about infants' abilities in this domain, and about how knowledge of simple facts in this domain emerges in development? Pursuing this question leads directly to puzzling patterns of evidence. These puzzles in turn point to theoretical challenges requiring, often enough, broadly philosophical solutions. Identify those puzzles and distinguish candidate solutions. In the best case, one of the candidate solution's predictions will turn out to be largely correct, and we will all have taken a tiny step towards understanding the developmental emergence of knowledge.

This is developmental philosophical psychology. It is not concerned with questions about the nature of developmental psychology as a science. Instead its aim is developmental psychology's aim: to explain the emergence, in development, of agentic, mental and social aspects of life. Its proponents, the philosophers, are under-labourers on a broadly scientific project. Their primary task is to be puzzled by discoveries, ideally in ways that will eventually promote further discoveries.

1.1. Two Breakthroughs

Can developmental philosophical psychology take us further than Plato got? Maybe. Two relatively recent scientific breakthroughs promise to shift thinking away from myths and closer to the minds and actions of actual humans. The first breakthrough concerns social interaction. It is the discovery that pre-verbal infants enjoy surprisingly rich social abilities. These may well facilitate the subsequent acquisition of linguistic abilities and enable the emergence of knowledge (as variously argued by several people, including for example Tomasello et al. 2005; Meltzoff 2007; Csibra and Gergely 2009).

A second breakthrough involves the use of increasingly sensitive—and sometimes controversial—methods to detect expectations without relying on subjects' abilities to talk or act. These methods have revealed that, from the early months of life on, infants have sophisticated abilities to track physical objects and their causal interactions, actions, mental states and more besides (e.g. Spelke 1990; Baillargeon, Scott and He 2010). The mental states underpinning these abilities surely play a role in the emergence of knowledge.

Although each of these breakthroughs has been extensively discussed, they are rarely considered together. There may be an opportunity to make progress by combining the breakthroughs. My guess is that development is like climate change in one respect. Lots of different mechanisms are simultaneously at work, and many interact with each other. To make progress we need to identify various mechanisms and understand their interactions. What follows is an attempt to show, by closely following what has been discovered so far, that understanding the emergence in development of knowledge will eventually require somehow bringing together the abilities that infants mani-

fest in the very first months of life concerning physical objects, minds and actions and their abilities to act jointly with those around them.

Before we get to the details, let me outline a little theoretical background.

1.2. Knowledge

The question we face—How do humans come to know about objects, actions, and minds?—is a question about knowledge. Answering this question depends on discovering when humans come to know what. And making these discoveries in turn depends on being able to distinguish really knowing something from merely manifesting some symptoms associated with knowledge.

Imagine an infant who seems to want a toy and, when given the chance, immediately searches for it in exactly the place it was lost. She is acting as if she knew where it was lost, so exhibiting a symptom of knowledge. But does she really know? Maybe not. If the state underlying her searching actions were locked to an arbitrarily limited range of actions, say, then it would not be knowledge. So what is distinctive of really knowing something?

In what follows, I take for granted that knowledge is constitutively linked to practical reasoning and to inference. Let me explain. Knowledge is the kind of thing that can typically influence how you act when you act purposively, and it is the kind of thing that can influence purposive actions in any domain at all. Knowledge is also the kind of thing that you can sometimes arrive at by inference, and which can enable you to make new inferences in any domain at all. A state that is not linked to practical reasoning and to inference in these ways is not knowledge.

I also take for granted that knowledge states are inferentially integrated with other attitudes like beliefs, desires and intentions. This does not mean, of course, that people are invariably rational. Instead the idea is this. One striking fact about many humans is that, at times, they achieve a kind of harmony in what they know, believe, intend, desire and do. Sometimes, some of their thoughts and actions come to be approximately rationally related. Now it may be that this is rare, or even highly unusual, in humans. But however infrequent, since it is not an accidental occurrence it stands in need of explanation. And the explanation, or part of it, involves processes of practical reasoning and inference. In saying that knowledge states are inferentially integrated with other attitudes like beliefs, desires and intentions, part of what I mean is that these they can (albeit perhaps rarely) come to be nonaccidentally related in ways that are approximately rational thanks to processes of inference and practical reasoning.

But there is more to being inferentially integrated. When humans are functioning at their best, they characteristically bring thoughts and actions into harmony unless something prevents them. This is the other part of what

I mean by saying that knowledge states, beliefs and the rest are inferentially integrated: in the absence of obstacles such as time pressure, distraction, motivations to be irrational, self-deception or exhaustion, approximately rational harmony will characteristically be maintained among currently active knowledge states, intentions and other attitudes.²

These facts about knowledge are almost too simple to mention. But they will turn out to be critical for distinguishing really knowing something from merely manifesting some symptoms associated with knowledge. The hypothesis that someone knows something generates the prediction that, in the absence of obstacles, she can manifest this knowledge in almost any situation.

If it is locked to an arbitrarily limited range of actions, or if it is used in response to an arbitrarily limited range of events, then it is not knowledge.

1.3. A Crude Picture of the Mind

Knowledge and the other attitudes contrast with perceptual representations. These are those postulated by scientific theories to explain processes such as edge detection or the computation of relative distances (see Palmer 1999, for an introduction).

Knowledge also contrasts with motor representation, which is less familiar but will be important later. Imagine being in a coffee shop where the servers use little round trays. You are watching the servers as they remove items from trays they are carrying around. As they lift a mug from a tray, the tray remains stable. How do the servers do this? They are not deliberating about the forces involved (or not usually). But nor is this a mindless physiological change. Instead it involves anticipation of the effects of their own actions, as you can discover by removing an item from the tray when a server is not looking—this can easily cause them to drop everything on the tray. Anticipatory control of action is one of things motor representations enable (see Rosenbaum 2010, for an introduction). They are those representations of actual, possible, imagined or observed actions and their effects which are characteristically involved in preparing, performing and monitoring sequences of small actions such as grasping, transporting and placing a mug.

Unlike knowledge states, perceptual and motor representations are plausibly not inferentially integrated with beliefs, desires, intentions and other attitudes. You can have perceptual experiences of the relative sizes, colours or locations of objects which are incompatible with what you know and believe. Such cases—illusions—are not due to you simply failing to make an

² I adapt the term 'inferential integration' from Stich's discussion of beliefs. According to him, for beliefs to be inferentially integrated is for there to be 'generally a huge number of inferential paths via which a given belief can lead to most any other' (Stich 1978, p. 506).

inference. Nor are they symptoms of self-deception or a divided mind. They are consequences of the fact that perceptual processes are, to an interesting extent, distinct from the inferential processes in which knowledge states feature. This is why there are illusions, and, more generally, why a single event can result in multiple incompatible representations.³

Let us take as our starting point a crude but quite standard picture of the adult mind. The mind comprises at least three kinds of states and processes:

1. epistemic (that is, knowledge-related),
2. motoric, and
3. perceptual.

The three kinds of process are to an interesting extent distinct from each other, and the three kinds of state are not inferentially integrated in the above sense.

When we explore recent discoveries about infants' abilities, we will see that they do not fit neatly with this crude picture of the mind. They appear to be in states which are not epistemic, not motoric and not perceptual. This will be a key theme in following chapters: understanding the developmental emergence of knowledge requires identifying states which do not fit neatly into the crude picture of the mind. One of the major unresolved challenges is finding a good way to revise the crude picture, one that can generate novel predictions.

1.4. Core Knowledge

The need to identify states which do not fit neatly into the crude picture of the mind has been discussed by Davidson, although on his view the need arises for philosophical reasons rather than as a consequence of any scientific discoveries. He writes:

'The difficulty in describing the emergence of mental phenomena is a conceptual problem [...] In [...] the evolution of thought in an individual, there is a stage at which there is no thought followed by a subsequent stage at which there is thought. To describe the emergence of thought would be to describe the process which

³ There is a further respect in which knowledge contrasts with perceptual and motor representations. Knowledge is a pretheoretical notion which features in social, legal and ethical contexts. By contrast, perceptual and motor representations are theoretical postulates. Their usefulness hinges on their roles in the best available scientific theories of perceiving and acting. In cognitive science, phenomena associated with knowledge are things to be explained whereas perceptual and motor representations are things which explain.

leads from the first to the second of these stages. What we lack is a satisfactory vocabulary for describing the intermediate steps' (Davidson 2001, p. 127).

Where Davidson uses the word 'thought', I am using 'knowledge'. This difference is unimportant here (because of inferential integration).

Davidson goes on to say that the problem cannot be solved:

'if you want to describe what is going on in the head of the child when it has a few words which it utters in appropriate situations, you will fail' (pp. 127–8).

But will we fail? Since describing 'what is going on in the head of the child' is the focus of much developmental psychology, perhaps there are some ideas that will help.

One key idea from developmental psychology is that of core knowledge (Spelke et al. 1992; Carey and Spelke 1996; Spelke 2000). As Carey puts it, the central claim is this:

'there is a *third type* of conceptual structure, dubbed "core knowledge" ... that differs systematically from both sensory/perceptual representation[s] ... and ... knowledge' (Carey 2009, p. 10; my emphasis).

Core knowledge states feature in core systems, which are 'largely innate, encapsulated, unchanging, arising from phylogenetically old systems, and built upon the output of innate perceptual analyzers' (Carey and Spelke 1996, p 520). For many domains of knowledge, the best-supported developmental theories postulate the existence of a core system as something more primitive than knowledge. These core systems are thought to provide a basis for the developmental emergence of knowledge.

I emphasised 'third type' in the quote just above because core knowledge is supposed to be a state distinct from knowledge (and belief). In terms of the above crude picture, which distinguishes epistemic, perceptual and motor representations,⁴ we should really say that core knowledge is supposed to be a *fourth type* of state.

A key question in what follows is whether Spelke, Carey and others are right that explaining development requires postulating a third or fourth type of mental state, something which is not epistemic, perceptual or motoric. This question is linked to a longstanding conflict between two stories about development.

⁴ Distinguishing perceptual from motor representations will matter in Chapters 7 and 10.

1.5. Two Stories

Any story has to accommodate the breakthrough discovery that infants, even in their first months, have sophisticated abilities to track objects, causal interactions, numerosity, actions, mental states and more besides in infants.

Perhaps the theoretically simplest (in a good sense) way to do this is to invoke knowledge. According to one story, infants' earliest abilities to engage with objects, actions and mental states are based on knowledge. From as early as they can manifest these abilities, they know some general principles in whatever sense the adults do. They use this initial knowledge of general principles together with perceptual information to make inferences about particular things—objects, actions, minds and the rest. Infants differ from adults only in that they do not know very much whereas, when things go well, adults know more. Development is essentially just a matter of acquiring more knowledge.

A strong case for this first story can be made, as we will see for one domain of knowledge in Chapter 3. In essence, hypothesising that infants know certain things enables us to characterise their abilities in a way that is both theoretically simple and mostly accurate. The problem is that such hypotheses also systematically generate incorrect predictions.

The other story is about core knowledge. On this story, infants' earliest abilities to engage with things are based on core knowledge of general principles (rather than on knowledge proper). This core knowledge somehow provides a basis for the acquisition of knowledge. And when infants do succeed in acquiring knowledge, this does not change their core knowledge. Instead their core knowledge is constant through life, even if it conflicts with things they later come to know. So in many domains there are two types of state that can influence cognition and behaviour, core knowledge and knowledge proper.

The second story aims to accommodate the breakthrough discovery about infants' early abilities while avoiding the incorrect predictions generated by conjectures about knowledge. Accordingly, the case for this second story is based on both evidence used to support the first story and evidence against it.

Although the second story does have the advantage of not generating incorrect predictions, it does have another weakness. As core knowledge is usually characterised, this story fails to generate relevant predictions (see Chapter 5).

At this point we reach an impasse. The challenge is to characterise what is going on in an infant's mind, and what is driving her actions, before she has any relevant knowledge states or beliefs. Neither of the standard stories seems adequate. One is simple and predictively strong but generates incorrect predictions; the other is theoretically complex but predictively weak. What to do?

1.6. Development Is Rediscovery

Here is a preview of my story. It starts from the idea that core knowledge is not one thing: it lacks both unity and uniformity. The phenomena associated with core knowledge are not exclusively either perceptual or motoric, although some are broadly perceptual and some are broadly motoric. And they comprise different kinds of representations and cognitive structures in different domains.

If this is right, current attempts to provide general theories of core knowledge are misguided because they assume unity (core knowledge is one thing) and uniformity (core knowledge is the same thing in different domains). Fortunately such theories are also mostly unnecessary. We can make good progress in understanding ‘core knowledge’ (or whatever else you would like to call infants’ earliest cognition of objects, actions, minds and the rest). This is because infants’ abilities and their limits can sometimes be explained by representations, structures and processes which have been identified and studied independently of developmental theories. As we will see, in some domains there is evidence to support detailed conjectures from which novel predictions flow (Chapters 6, 7, 10 and 14).

Thinking about core knowledge in this way motivates construing development as a process of rediscovery. Infants’ sophisticated abilities to track objects, actions and mental states reveal that, even in the first months of life, their cognition and action is influenced by an impressive range of truths and useful falsehoods about the natures of these things. (This is the breakthrough discovery that led to theories about core knowledge, of course.) But these truths and falsehoods, if represented at all, are represented in ways that are cut off from knowledge. The representations are not inferentially integrated with knowledge and cannot lead to the acquisition of knowledge by inference. Further, their causal interactions with knowledge may involve intermediaries which either lack intentional features entirely or else have intentional features that are only very distantly related to those of the cognitive structures they link. That is, they may be intentionally isolated from knowledge.

Imagine a single organization with different parts. One part has information about criminals which enables it to predict and detect break-ins. This is the security department. Another part of the organization lacks this information. But instead of communicating internally, this part goes out and learns everything from scratch; and what it learns is never fed into the security department. This is the research department. From the point of view of this department, it is almost as if the organization starts with no knowledge of criminality at all. But only almost—for the learning done in the research department depends on the operations of the security department to keep it safe, and to flag up criminal events. This is rediscovery: one part goes out and learns afresh about things related to those already relied upon by another

part of the organization; meanwhile the other part's ongoing contribution to the organization's flourishing continues unaffected by any new learning.

You might think this analogy shows that the idea of rediscovery is flawed. After all, this kind of organizational behaviour is routinely condemned as failure. But primate minds are special kinds of organization. They enable rapid learning in ways which are open to radical revision and can therefore yield radically mistaken conclusions. This learning should not be contaminated by assumptions which happen to enable the animal to function from the first months of life onwards, and which are effective even when it has little experience of the world and limited social interactions. And nor should the mistakes it makes in learning compromise its ability to engage effectively with physical and social aspects of the world. Both requirements can be met thanks to the inferential and intentional isolation of the security department from the research department—or of the representations involved in early-developing abilities from knowledge.

If inferential and intentional isolation mean that development is rediscovery, we face a challenge. How could anything like core knowledge ever facilitate the developmental emergence of knowledge?

The phenomena associated with core knowledge and the infant's knowledge and beliefs can be connected indirectly, via the her behaviour, experience and attention. To illustrate with a comparatively simple example, consider the development of face perception. From birth there is a mechanism present in humans (and chickens) that uses crude heuristics to identify faces and generates orienting reflexes. There is also a later-developing mechanism which uses more sophisticated, learned principles geared to features of conspecifics and enables smooth tracking of moving faces. Crucially the two mechanisms do not share any representational resources; the second mechanism does not rely on signals carrying information about whether the first mechanism has detected a face. Instead they are linked only indirectly, via behaviour: the first orients the baby's eyes to faces, thereby indirectly providing inputs necessary for learning in the later-developing mechanism (Johnson and Morton 1991; Haan 2002). This is one illustration of how, despite their intentional isolation, an early-developing ability can facilitate the emergence of a later-developing one indirectly, by influencing behaviour, experience and attention.

The case of face detection also shows, by the way, that arrangements involving multiple mechanisms are not necessarily wasteful. In this case the arrangement makes sense given the different needs humans have at different times of life. The first mechanism relies on fixed heuristics that work well in the particular position newborn infants, who cannot support their heads, mostly find themselves in. This enables infants to reliably identify, and orient to, faces from birth. Orienting to faces provides useful input for the second, later-developing and more flexible mechanism. The first does not tell the second which things are faces or how to detect them: they are intentionally

isolated. And because it is not constrained by the fixed heuristics of the first mechanism, the second mechanism eventually enables older infants to detect faces in a wider range of situations. Rediscovery costs time and cognitive effort, of course. But it provides benefits in accuracy and flexibility.

One crucial feature is missing from this illustration. If development is rediscovery, and if rediscovery is something achieved by acting together with those who care for us, then infants' social skills are indispensable drivers of knowledge. That is why the first part of this book focusses on infants' abilities concerning physical objects and the second on their social skills. I will not be attempting to construct an account of how knowledge emerges in development; but I will introduce the two breakthroughs that, eventually, must somehow be combined in explaining the developmental origins of knowledge.