BETWEEN
PERCEPTION AND ACTION

BENCE NANAY

Between Perception and Action

Between Perception and Action

Bence Nanay

To Felicitas

Contents

Acknowledgements

This book is about perception, action, and what's in between. It's about the whole mind, in other words. It is written for philosophers, but not only for philosophers: it is also for psychologists or other empirical researchers of the mind with theoretical interests, as well as people who are just interested in understanding how the human mind works and how it is very similar to animal minds.

I started exploring the central claims of this book almost 20 years ago. As far as my philosophical attention deficit disorder allowed, I have been pursuing these claims ever since. In 1995, I published a very bad paper in a psychology journal on how we should reject Gibson's views about perception while preserving the insight about the action-oriented nature of perception. Then I wrote a still pretty bad MPhil thesis at Cambridge University on the "special" mental states that mediate between sensory input and motor output, and then an only slightly better PhD at the University of California, Berkeley on action-oriented perception. I am grateful to my various teachers and colleagues throughout these years (and admire their patience) for not discouraging me from pursuing this project, given how unconvincing some of my arguments really were. The hope is that the arguments in this book are much more convincing.

Given that I have been talking to philosophers, psychologists, ethologists, and neuroscientists about the topic of this book for nearly 20 years, there is no way I can remember everyone who helped me with thinking more clearly about the topic of the book, or even provided written comments on earlier written-up versions of some of the ideas in the book. So I won't even try—if I tried, I would, no doubt, leave out many people. But I am especially grateful to those who read the whole penultimate version of the manuscript and provided detailed feedback: Jonathan Cohen, two anonymous referees from Oxford University Press, and my postdocs, Maja Spener, Carolyn Dicey Jennings, Will Davies, and Craig French.

Various parts of the book were presented at the following symposia, conferences, and workshops: American Philosophical Association, Pacific Division (San Diego), Annual Conference of the Cognitive Science Society (Vancouver), American Philosophical Association,

Pacific Division (San Francisco), American Philosophical Association, Pacific Division (Vancouver), Annual Conference of the Cognitive Science Society (Berlin), Intuitions, Experiments and Philosophy Conference (Nottingham), Social Cognition Workshop (Cambridge), Evolution, Non-propositional Intentionality Workshop (Geneva), Co-operation and Rationality Conference (Bristol), Philosophy of Psychology, Neuroscience and Biology Conference (Oxford), Empathy/ Sympathy Conference (Antwerp), Social Science Roundtable (Paris), Systematicity Workshop (San Jose), Phenomenal Qualities Conference (Hertfordshire), Animal Cognition Conference (Bochum), Philosophical Insights Conference (London), Philosophy of Science Association Biannual Meeting (San Diego), 5th Online Consciousness Conference, New Waves in Philosophy of Mind Online Conference.

I also gave talks based on the material in this book at the University of Bochum, University of Fribourg, Institut Jean Nicod, University of Edinburgh, Simon Fraser University, University of British Columbia, University of Cardiff, University of Turin, University of Geneva, Rice University, Syracuse University, York University, Concordia University, and the University of California, Berkeley. I am extremely grateful for the comments and grilling I received after these talks. I also learned a lot from teaching some of the material in the book—both at PhD seminars and at the undergraduate level—at Syracuse University, the University of British Columbia, and the University of Antwerp.

Bence Nanay

Cambridge
February 19, 2013

1

Introduction

We are only philosophers when at work in our study: so soon as we indulge freely in the impressions of our senses, we are children like everyone else. Abstract speculation is a forced state of mind: our natural condition is that of everyman.

Jean-François Marmontel: *Poétique Françoise*, 1763.

1.1 The icing and the cake

There are two general approaches towards trying to understand the human mind. The first one is to take our sophisticated, highly intellectual, rational, linguistic, and uniquely human mental capacities to be paradigmatic, and try to understand how our minds work by focusing on them.

The second way of trying to understand the human mind emphasizes the continuity between human and animal minds. It therefore focuses on those simpler mental capacities that non-human animals also possess: those mental processes that make it possible for us to perform actions successfully—to put food in our mouth, or to get around without bumping into things.

As philosophers are in the business of giving rational, sophisticated, highly intellectual arguments, a tempting route for philosophers of mind is to project their own methodology onto the phenomenon they are trying to explain, and consider the human mind as primarily consisting of rational, highly intellectual, sophisticated mental capacities. I think we should resist this temptation.

My aim is to argue that the vast majority of what goes on in our mind is very similar to the simple mental processes of animals. Our complex, sophisticated, rational, and linguistic abilities could be described as the icing on the cake. Although the icing is what makes the cake look great,

we should try to understand the cake without worrying too much about the icing. In other words, the right methodology for philosophy of mind is to understand those simple mental capacities that we share with animals first, and then explain those uniquely human, highly intellectual mental capacities that make the human mind so remarkable. It is important that I am not suggesting that the human mind is to be understood as the animal mind plus some extra features—imagine instead a cake where the icing nicely seeps into the dough. But if we want to understand the human mind, we need to start with the simplest mental capacities and work our way towards the more complex ones.

To stretch the cake analogy even further, the aim of this book is to identify what could be described as the main ingredients of this cake. The basic units of our linguistic, higher order mental processes are conceptually/linguistically structured propositional attitudes: beliefs, desires, thoughts. But we have no reason to suppose that the basic units of those simple mental capacities that we share with animals would also be such propositional attitudes. In fact, the point could be made that when we describe the human mind as consisting of beliefs and desires, we are mirroring language. The basic units of language are sentences or utterances that express propositions; thus, it would be tempting to say that the basic units of our mind must be mental states that also express propositions: propositional attitudes. But if we maintain that our under- standing of the human mind is not to be modeled on language or even on uniquely human linguistic thoughts, then we have no reason to accept that beliefs and desires are the basic building blocks of the human mind.

The need to look for new candidates to replace beliefs and desires as the basic building blocks of the human mind is not new. Gareth Evans, for example, suggested that we should consider information-carrying states to be the basic units instead (Evans 1982). I would like to make a more specific proposal. My claim is that the vast majority of what is going on in the human mind can be understood if we consider what I call "pragmatic representations" to be the basic units of our mental capacities.

The human mind, like the minds of non-human animals, has been selected for allowing us to perform actions successfully. And the vast majority of our actions, like the actions of non-human animals, could not be performed without perceptual guidance. My claim is that the mental state that mediates between sensory input and motor output is the basic

building block of the human mind. I call mental states of this kind pragmatic representations.

1.2 Pragmatic representations

What mediates between sensory input and motor output? This is probably the most basic question one can ask about the mind. There is stimulation on your retina, something happens in your skull, and then your hand reaches out to grab the apple in front of you. What is it that happens in between? What representations make it possible for you to grab this apple? My answer to these questions is that it is pragmatic representations that mediate between sensory input and motor output; it is a pragmatic representation that makes it possible for you to grab the apple.

Pragmatic representations are, at first approximation, the representational components of the immediate mental antecedents of action. They are also genuine perceptual states. The immediate mental antecedents of action are what make actions genuine actions. They constitute the difference between actions and mere bodily movements. They guide our ongoing bodily activities. And pragmatic representations are the representational components of these immediate mental antecedents of action.

These mental states represent the world as being a certain way: they are *about* the world, they *refer to* the world. In other words, they have representational *content*: they represent objects as having certain properties. This, however, does not mean that they must have a syntactically articulated propositional structure, or that they really are sentences written in some mental language. Pragmatic representations can be correct or incorrect. If they are correct, they guide our bodily activities well. If they are incorrect, they guide us badly.

What properties do pragmatic representations represent objects as having? Suppose that you want to pick up a cup. In order to perform this action, you need to represent the cup as having a certain spatial location, otherwise you would have no idea which direction to reach out towards. You also need to represent it as having a certain size, otherwise you could not approach it with the appropriate grip size. And you also need to represent it as having a certain weight, otherwise you would not

know what force you need to exert when lifting it. I call these properties "action-properties": action-properties are properties that need to be represented in order for the agent to perform the action. Pragmatic representations attribute action-properties: they represent objects in an action-oriented manner (Nanay 2011a, 2011b, 2012a). And they typically attribute these action-properties unconsciously.

As pragmatic representations are supposed to play an important role in action, and as they are supposed to be perceptual states, if we accept that pragmatic representations are crucial for understanding our mental life then this will have radical consequences both for debates in philosophy of perception and for debates in philosophy of action.

Let us take philosophy of perception first. Pragmatic representations are *bona fide* perceptual states. It is an important question in the philosophy of perception what properties we perceive objects as having. Shape, size, and color are obvious candidates. It is much less clear whether we represent sortal properties perceptually: whether we literally see objects as chairs and tables and not just infer that they are. I argue that we sometimes also perceive objects as having action-properties.

This should not sound particularly surprising. Our (extremely complex) perceptual system was selected for helping us to perform actions on which our survival depended. It is hardly surprising, then, that it was selected for representing objects as having properties that are relevant to the performance of our actions: as having action-properties. As Tyler Burge says, "since perception guides action, it is not surprising that perceptual kinds mesh with [our] activities" (Burge 2010, p. 324).

Importantly, my claim is that we perceive objects as having action-properties *some of the time*, not that we always do so. Pragmatic representations are perceptual states, but not all perceptual states are pragmatic representations. Very often we perceive objects in a way that is not action-oriented—for example, when we are admiring the landscape, sitting on a bench, in a detached manner, without any urge to perform any action at all.

Now let us see the relevance of pragmatic representations for the philosophy of action. One of the most important questions of philosophy of action is this: what makes actions actions? How do actions differ from mere bodily movements? The difference is some kind of mental state that triggers, guides, or maybe accompanies the bodily movement. But what is this mental state?

I argue that these immediate mental antecedents of our actions are partly constituted by pragmatic representations. These pragmatic representations attribute properties, the representation of which is necessary for the performance of the action, and they do so perceptually: they perceptually guide our bodily movement. But if this is true, then the belief–desire model of motivation needs to be adjusted or maybe even discarded.

According to the "belief–desire model" of motivation, beliefs and desires are necessary for motivating us to act. Perception may play a role in bringing about the belief that our action is based on, but it does not play any more substantial role. Here is an example. I look out of the window and I see that it is raining outside. I form a belief that it is raining outside. I have a desire not to get wet and this, together with my further (instrumental) belief that the best way of not getting wet in rain is to take an umbrella, leads to the forming of an intention to take an umbrella, and this intention triggers my action of taking my umbrella (see, e.g., Smith 1987, Davidson 1980, the origins of the account may go back to David Hume). In this model, beliefs and desires mediate between sensory input and motor output. Although this model does describe the way in which some of our (highly intellectual and arguably uniquely human) actions come about, it only applies to a very small minority of the actions we perform. Some of our perceptual states, namely pragmatic representations, play a more crucial role in bringing about our actions than has been acknowledged.

In the case of most of the actions we perform, the only mental states (more precisely: the only representational states) that mediate between sensory input and motor output are pragmatic representations—and not a set of beliefs and desires, as the classical belief–desire model suggests. Further, even in those cases when beliefs and desires play an essential part in bringing about our actions, pragmatic representations still need to be involved (for example, when I actually pick up the umbrella).

We have seen two aspects of the philosophical relevance of pragmatic representations: each time we are performing an action, we have a pragmatic representation. And many of our perceptual states are pragmatic representations.

But there are further reasons to be interested in mental states of this kind. Animals and small children are capable of performing goal-directed actions, such as running away from predators or chasing prey.

But if they are, they must be able to have pragmatic representations. Hence, even organisms that may be incapable of entertaining complex thoughts and beliefs must be able to have pragmatic representations.

These mental states are both phylogenetically and ontogenetically quite basic. It can be argued that the first mental representations that appeared in the course of evolution were pragmatic representations. And, similarly, the first mental representations that appear in the course of child development are also likely to be pragmatic representations. If we want to explain the minds of non-human animals or small children, we would be well advised to focus on mental states of this kind.

The general rhetoric I am following here is to examine the simple mental processes that animals and humans are equally capable of, and postpone the analysis of uniquely human, complex, sophisticated mental capacities. But if we accept the general framework I propose, we can go a bit further and at least begin to explain some of our more complex mental capacities, such as our engagement with, and understanding of, others, and even basic emotions such as empathy.

As we have seen, pragmatic representations represent objects in a way that is relevant to *my* action. But there is an intriguing phenomenon in the vicinity: mental states that represent objects in a way that is relevant to *someone else's* action. I argue that these mental states constitute our simplest and most basic capacity to engage with others. I call this way of engaging with others *vicarious perception*, and argue that a number of big debates—not only in philosophy, but also in cognitive ethology and developmental psychology—about our ability to attribute mental states to others (an ability usually called "theory of mind") could be resolved if we restrict their scope to vicarious perception.

In the last three decades, the concept of "theory of mind" has been at the center of interest in philosophy of mind, psychology, and primatology. Some important questions about "theory of mind" are the following: Do non-human animals have theory of mind? How does theory of mind develop in ontogeny? What mental processes make theory of mind possible in humans? What are the neural underpinnings of theory of mind?

While these questions are difficult to tackle (and there has been no sign of any consensual answer) so long as they are about "theory of mind," if we take them to be about vicarious perception (and not theory of mind) we get straightforward and nontrivial answers. More precisely,

it can be argued that all experiments that are supposed to show that non-human primates have theory of mind in fact demonstrate that they are capable of vicarious perception. The same goes for the experiments with infants less than 12 months old. If we shift the emphasis from theory of mind to vicarious perception, we can make real advances in understanding the origins of social cognition.

1.3 Product differentiation

The traditional way of thinking about the human mind has been to explain it using the analogy of symbol-manipulating serial computers. Sensory input comes in, it gets coded into some kind of symbolic, sentence-like propositional format and is matched against other symbols, and this sometimes gives rise to motor output (Fodor 1981, 1983, Davidson 1980). This is the picture the traditional belief–desire model uses, where propositional attitudes (beliefs, desires, intentions) mediate between the sensory input and the motor output. There is no agreement about just what kind of propositional attitudes these mediating states are—and it very much depends on one's concept of propositions. What I take to be the traditional way of thinking about the human mind takes the mediators between sensory input and motor output to be propositional attitudes whose content is syntactically, conceptually, and maybe even linguistically articulated. I will call this view the "computationalist/ propositionalist" account (see Figure 1). Those who take propositions to be necessarily syntactically, conceptually, and maybe even linguistically articulated could just read "propositionalist" here. I add the "computationalist" label because those who have a weaker notion of propositionality may object to calling this view propositionalist per se. Thus the label: "computationalism/propositionalism."
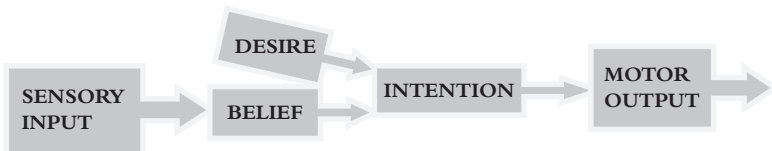
Figure 1 Computationalism/propositionalism

The most important alternative to the computationalist/propositionalist account of the human mind could be labeled the anti-representationalist/enactivist approach. According to the enactivist, perception is the active exploration of one's environment, which can be described without talking about representations at all (Noë 2004, Hurley 1998, Gibson 1979). The proponents of this approach reject over-intellectualizing the mind by shifting the emphasis from mental to bodily activities: by emphasizing the importance of bodily coping—skills and abilities that are not really (or not exclusively) mental—they tend to emphasize the embodied nature of cognition. And this tends to amount to denying that perception is supposed to be described in representational terms at all.

According to this anti-representationalist/enactivist view, "perception is not a process of constructing internal representations" (Noë 2004, p. 178, see also Ballard 1996, O'Regan 1992). The main enactivist claim is that we have all the information we need in order to get around the world out there, in the world. So we do not need to construct representations at all and, more specifically, we do not need perceptual representations. As Dana Ballard put it, "the world is the repository of the information needed to act. With respect to the observer, it is stored 'out there', and by implication not represented internally in some mental state that exists separately from the stimulus" (Ballard 1996, p. 111, see also Brooks 1991, Ramsey 2007, Chemero 2009, Hutto and Myin 2013). Sensory input and motor output are so closely intertwined in a dynamic process that we do not need to posit any representations that would mediate between them (see Figure 2).
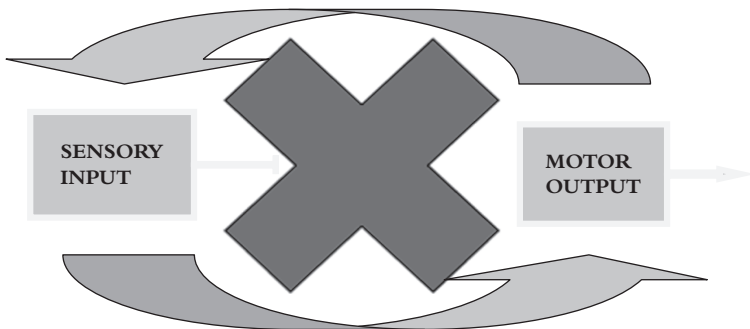
Figure 2 Anti-representationalism/enactivism

I think that this is a mistake, and that rejecting talk of mental representations altogether because of the justified mistrust of a specific kind of representation—namely, conceptually/linguistically structured propositional attitudes—is a bit like pouring out the baby with the bathwater. These kinds of propositional attitudes are not the only kind of representations. My view is that the mind is to be understood in terms of representations, but that these representations are not conceptually or linguistically structured, nor are they uniquely human. They are better compared to the mental representations of the predator that make it possible for it to catch its prey. These representations are perceptual representations and inherently action-oriented (see Figure 3).

| SENSORY INPUT | → | PRAGMATIC REPRESENTATION | → | MOTOR OUTPUT | → |

Figure 3 Pragmatic representations

In short, the novelty of my proposal is to carve out an intermediary position between the enactive/antirepresentationalist and the computationalist/propositionalist accounts. This intermediary position is supposed to inherit the explanatory advantages of both extremes.

It inherits the general view that representations are the basic building blocks of our mind from the computationalist/propositionalist accounts. The general upshot is simple: we should talk about the human mind in terms of representations. So the general theoretical framework is the same as the one the computationalist uses: sensory inputs give rise to some representations, which give rise to motor outputs.

But these intermediary representations in my framework and those used in the computationalist framework are very different. While the computationalist takes these representations to be propositional attitudes (that follow, more or less, a belief–desire model of connecting sensory input to motor output), I argue instead that these representations are (a) perceptual representations and (b) inherently action-oriented.

The main advantage of my approach over the computationalist one is that I can talk about the mental representations of animals and small children (as well as the mental representations of adult humans that make most of our everyday actions possible) without

over-intellectualizing what goes on in their mind. In other words, my framework can capture the anti-cognitivist and anti-computationalist intuitions that fuel the enactive approach without having to give up the concept of representation altogether.

An additional, very important, advantage of my approach is that if we consider the representations that mediate between sensory input and motor output to be pragmatic representations, then we can have a better (and richer, more complex) understanding both of perception and of action. One way of summarizing my approach would be to say that perception and action take up more of what goes on in our mind than the computationalist suggests. Pragmatic representations could be thought of as the overlap between perception and action: they are genuine perceptual states and they are part of what makes actions actions. It is in this sense that they can be said to provide direct mediation between sensory input and motor output.

So far, I have outlined the main difference between my view and the computationalist position. The difference between my view and the enactivist position is even more fundamental: they reject the concept of representation, whereas I am endorsing it. What are the explanatory advantages of my view over the enactivist approach? I argue elsewhere that some important empirical phenomena about the human mind, namely action-guiding vision and the multimodality of perception, cannot be sufficiently explained by the anti-representationalist approach (Nanay forthcoming c). More generally, if we are looking for an alternative to the over-intellectualized, language-based description of the mind, we need to give an account that has similar explanatory power. And I am doubtful that we can explain much about the mind unless we consider it to be a mechanism for representing the world. The important question is *how* it represents the world (not whether it does). My aim is to answer this question by analyzing the properties that pragmatic representations attribute to the perceived objects.

Another important difference between my view and enactivism is that although they both emphasize the intricate connection between perception and action, they do so very differently. In recent years, there has been a flood of theories that emphasize the interdependence between perception and action. My account could be thought of as a variant of these, so I need to explain how and why it differs from the existing theories of the interdependence of perception and action.

One recently popular way of cashing out this interdependence is to argue that perception is an active process (Lewis 1929, Noë 2002, 2003, 2004, Hurley 1998, Rowlands 2006). The idea is to think of perception *as* action, and urge to replace vision as the paradigmatic case of our perceptual processes with tactile perception, where the act of actively exploring the environment is built into the very notion of perceiving. A slightly different version of emphasizing the perception–action inter-dependence focuses on the bodily movements that allow us to perceive objects, or to perceive them better. We move our eyes and head, lean forward, and cup our ears in order to achieve the most efficient way of receiving sensory impressions from the perceived object.

Instead of interpreting perception *as* action, I claim that perception is (some of the time, not always) *about* (or for) action. My main interest is not to give a new description of the nature of perception, but to clarify how our action relies on our perception and how our perception carves up the world accordingly. In short, for me the main question is about what mediates between sensory input and motor output. And my answer is: pragmatic representations.

Finally, some aspects of my view may have reminded the reader of J. J. Gibson's theory of perception, or, more generally, of Gibsonian, or ecological, theories of perception (Gibson 1979, Turvey 1992, Chemero 2001, 2009—for a recent philosophical revival, see Prosser 2011, Campbell 2011). I need to be clear about where my view differs from this approach. Gibson is one of the most important proponents of the anti-representationalist tendency that I tried to differentiate my account from above. But Gibson's more specific theory of affordances may, on the face of it, seem very similar to the account I am defending here. So it is useful to highlight the main difference between Gibson's theory and mine.

Gibson emphasizes the importance of perceiving affordances, whereas one way of cashing out my claim is that we sometimes perceive objects as being relevant to our actions. Note the difference. Gibson denies that we perceive objects and attribute properties to them. For him, what we perceive are affordances, not objects—this claim is the source of some of the most influential philosophical objections to his account (see, for example, Fodor and Pylyshyn 1981). My view is much less revisionary. I do not urge any radical change in the way we think of the metaphysics of perception, like Gibson did. All I claim is that among

the numerous properties we attribute to the perceived objects, some are action-properties: properties the representation of which is necessary for the performance of one's action.

The structure of the book is simple: after introducing the concept of pragmatic representations, and arguing that they are perceptual states and that they play an important role in our mental life (Chapter 2), I examine the implications of this view for philosophy of perception (Chapter 3) and for philosophy of action (Chapter 4). In the last two chapters, I analyze some important variants of pragmatic representations: "pragmatic mental imagery"—that is, mental imagery that represents (the imagined) objects as having action-properties, as being relevant to our actions (Chapter 5); and "vicarious perception"—the phenomenon of perceiving things as being relevant not to one's own action, but to someone else's (Chapter 6).

2

Pragmatic Representations

Although the world we experience and the world in which we act appear to us to be one and the same, they really resemble the water above and below a mill, they're connected by a sort of reservoir of consciousness with a sluice-gate to regulate the level, pressure, and so forth, on which the influx and efflux depend.

Robert Musil, *Der Mann ohne Eigenschaften*, 1930.

2.1 Mediating between sensory input and motor output

Pragmatic representations are the representations that mediate between sensory input and motor output. The aim of this chapter is to argue that they are perceptual states, and that they play a very important role in our mental economy.

More precisely, the aim is to argue that there are mental representations that are both perceptual states and (part of what constitutes) the immediate mental antecedents of actions. My strategy will be to start out with the immediate mental antecedents of actions, and show of them that they are (very often) perceptual states. But first, I need to clarify what I mean by perception and what I mean by action.

I will have a lot to say about how to draw the line between actions and non-actions in Chapter 4. For now, it is enough to clarify that I will mainly focus on basic actions: actions that are not performed by performing another action (Danto 1963, 1965, Brand 1968). Raising my arm is a basic action. If by raising my arm I am also bidding in an auction, bidding is a non-basic action: I perform it by performing the action of

raising my arm. While a lot of debates in the last decades of philosophy of action have focused on non-basic actions, the starting point of my discussion is basic actions.

The account I propose for basic actions can be extended to non-basic actions easily. The action of switching on the light and the action of scaring away the burglars by switching on the light have the same immediate mental antecedent—the mental state that guides my hand to flip the switch. These two actions are very different in a number of respects, of course—for example, they may be motivated by very different beliefs/desires/prior intentions, or they may have very different consequences. Nonetheless, their immediate mental antecedents are the same. More generally, the immediate mental antecedent of a non-basic action is the same as the immediate mental antecedent of the basic action by means of which this non-basic action is performed. I will mainly use basic actions as examples, but what I say about them can be generalized to non-basic actions this way.

I also need to clarify what I mean by perception. First, perception can be conscious or unconscious. The discussion of perception is too often hijacked by those who are primarily interested in consciousness and conscious perception. I am interested in perception per se—whether or not it is conscious. In fact, as we shall see, pragmatic representations are typically unconscious, by which from now on I mean that they are not typically accessible to introspection.

Further, perception is not the same as sensory stimulation. Sensory stimulation is the starting point of perception—in the case of vision, retinal stimulation is the starting point of visual perception. But it is not the endpoint. There is a lot of visual processing after the retina and, more generally, there is a lot of perceptual processing after the sensory stimulation—some bottom-up, some top-down. I will discuss at length in Chapter 3 what the endpoint of this perceptual processing is: when perceptual processing gives rise to non-perceptual processing. At this point it is enough to make it clear that perception goes further than sensory stimulation—it remains to be seen just how much further it goes.

I will mainly use examples that involve vision. But everything I say about pragmatic representations applies equally to all the other sense modalities. I say more about the other sense modalities, as well as the multimodal nature of perception, in Chapter 3.

2.2 The immediate mental antecedents of actions

One of the classic questions of philosophy of action (or maybe even the most "fundamental question," see Bach 1978, Brand 1979) is the following: What makes actions actions? How do actions differ from mere bodily movements? What is the difference between performing the action of raising my hand and having the bodily movement of my hand going up? In short, what makes actions more than just bodily movements? Given that the bodily movement in these two cases is the same, whatever it is that makes the difference, it seems to be a plausible assumption that what makes actions actions is a mental state that triggers, guides, or maybe accompanies the bodily movements. If bodily movements are triggered (or guided, or accompanied) by mental states of a certain kind, they qualify as actions. If they are not, they are mere bodily movements.[1]

The big question is of course: What are the mental states that trigger (or guide, or accompany) actions? And there is no consensus about what these mental antecedents of actions are supposed to be. Myles Brand called mental states of this kind "immediate intentions" (Brand 1984), Kent Bach "executive representations" (Bach 1978), John Searle "intentions-in-action" (Searle 1983), Ruth Millikan "goal state representation" (Millikan 2004, Chapter 16), Marc Jeannerod "representation of goals for actions" or "visuomotor representations" (Jeannerod 1994, section 5, Jeannerod 1997, Jacob and Jeannerod 2003, pp. 202–4). I myself called them "action-oriented perceptual states" (Nanay 2012a) or "action-guiding perceptual representations" (Nanay 2011a).[2]

Note that all of these proposed mental states are representational states. More precisely, they seem to be representational states that

---

[1] Theories of "agent causation" deny this, and claim that what distinguishes actions and bodily movements is that the former are caused by the agent herself (and not a specific mental state of her). I leave these accounts aside because of the various criticisms of the very idea of agent causation (see Pereboom 2004 for a summary).

[2] This list is supposed to be representative, not complete. Another important concept that may also be listed here is John Perry's concept of "belief-how" (Israel et al. 1993, Perry 2001, see also Grush 2004, Hommel et al. 2001, Norman 2002). Also, there are important differences between these accounts of what makes actions actions—I will say more about these in Chapter 3 (see especially Section 3.2).

attribute properties, the representation of which is relevant or maybe even necessary for the performance of the action. And these representations guide our bodily movements. But I do not want to make the assumption that these mental states are representational states. All I assume is that they have a representational component.

To maintain neutrality, then, I use the term "the immediate mental antecedent of actions" as a place-holder for the mental state that makes actions actions: this mental state is present when our bodily movement counts as action, but is absent in the case of reflexes and other mere bodily movements. Thus, we can talk about the "immediate mental antecedents of actions" in the case of all actions. Intentional actions have immediate mental antecedents, but so do non-intentional actions. And autonomous intentional actions have immediate mental antecedents as much as non-autonomous actions (see Velleman 2000, Hornsby 2004).

As immediate mental antecedents of action are what makes actions actions, understanding the nature of these mental states is a task for philosophers of action that is logically prior to all other questions in action theory. In order to even set out to answer questions like "What makes actions intentional?" or "What makes actions autonomous?" one needs to have an answer to the question "What makes actions actions?" And the way to answer this question is to describe the immediate mental antecedents of action.

Many philosophers of action make a distinction between two different components of the immediate mental antecedent of actions. Kent Bach differentiates "receptive representations" and "effective representations" that together make up "executive representations," which is his label for the immediate mental antecedent of action (Bach 1978, see especially p. 366). Myles Brand talks about the cognitive and the conative components of "immediate intentions," as he calls the immediate mental antecedent of action (Brand 1984, p. 45). Leaving the specifics of these accounts behind, the general insight is that the immediate mental antecedent of action has two distinct components: one that represents the world (or the immediate goal of the action—more on the difference between the two in Chapter 3) in a certain way, and one that moves us to act. These two components can come apart, but the immediate mental antecedent of actions consists of both (at least in most cases—I will come back to the question about whether both of

these components are strictly necessary, as Bach and Brand assume, in Chapter 4).[3]

This book is about the representational component of the immediate mental antecedent of actions. Some may worry about the concept of "immediate" in this characterization, and about how one can or should individuate mental states that precede the action—especially if, as we have seen, this mental state is supposed to have two separate components (that may not be perfectly aligned temporally). In order to sidestep these issues, those who share this worry could read "the immediate representational antecedent of action" instead of "the representational component of the immediate mental antecedent of action" in what follows—this may even be a more straightforward way of making sense of pragmatic representations. I will keep on using the more complicated "representational component of the immediate mental antecedent of action" phrase, because the distinctions between the two different components of the immediate mental antecedent of action will play an important role in Chapter 4.

I call the representational component of the immediate mental antecedent of actions (or, alternatively, the immediate representational antecedent of actions) "pragmatic representations" (Nanay 2013, forthcoming c). Thus, it is true by definition that in order to perform an action we must have a pragmatic representation. But having a pragmatic representation does not necessarily manifest in an action, as it is the conative component of the mental antecedent of actions that moves us to act, and if we have the representational, but not the conative, component of the immediate mental antecedent of action, then the action is not performed.

A quick note about the scope of these claims, and the scope of the main claim of this chapter—namely, that pragmatic representations are perceptual states. When I say that all actions are triggered or accompanied by a mental state I call the "immediate mental antecedent of action," this also means that it is also true of mental actions. Some of our actions are impossible to perform without overt bodily movement. Some other actions can be performed without overt bodily movement—these actions

---

[3] This distinction between the cognitive and the conative components of the immediate mental antecedents of action is not entirely uncontroversial. I will discuss potential reasons against it in Section 2.3, as well as in Chapter 4.

are sometimes called "mental actions." Counting, multiplying, deciding, imagining, remembering, and reciting poetry to oneself are mental actions (Peacocke 2007). Because the emphasis of this book is on simple mental processes that we share with animals and small children, the focus of my discussion will be non-mental actions. But mental actions also have immediate mental antecedents. And these immediate mental antecedents also have a representational component—in order to visualize an apple successfully I need to represent how apples look, and in order to recite a poem to myself I need to know the lines. But I don't think we have any reason to believe that the representational components of the immediate mental antecedents of mental actions are perceptual states (although some may be quasi-perceptual states, such as mental imagery, see Shepard and Metzler 1971). The argument I will present for the claim that pragmatic representations are perceptual states only applies to non-mental actions. So when I talk about actions in what follows, I will assume that they are non-mental actions. And when I talk about pragmatic representations, I will assume that they are the representational components of the immediate mental antecedent of a non-mental action. I will say more about how my account could be generalized to at least some mental actions, especially imagining, in Chapter 5.

Pragmatic representations are genuine mental representations: they represent objects as having a number of properties that are relevant to performing the action. Let us go back to the example of picking up a cup. In order to perform this action, your pragmatic representation needs to represent the cup as having various properties. It needs to represent it as having a certain spatial location, otherwise you would have no idea which direction to reach out towards. It also needs to represent it as having a certain size, otherwise you could not approach it with the appropriate grip size. And it also needs to represent it as having a certain weight, otherwise you would not know what force you need to exert when lifting it.

Pragmatic representations can be correct or incorrect—any of these properties can be correctly or incorrectly represented. If they are correct, they are more likely to guide our actions well; if they are incorrect, they are more likely to guide our actions badly. And, as we shall see, they typically attribute these properties unconsciously.

My aim in this chapter is to argue that they are perceptual representations.

2.3 The direction of fit

I said that pragmatic representations are genuine representations. If this is so, then we need to clarify their "direction of fit" (Anscombe 1957, Platts 1979, Searle 1983, Smith 1987, Humberstone 1992). Some representations represent the world as being a certain way. They attribute properties to objects. If the represented objects have these properties, the representation is correct. If they do not have these properties, it is incorrect. Representations of this kind have a "mind to world" direction of fit.

Some other representations, in contrast, have a "world to mind" direction of fit: they do not describe how the world is, but prescribe how the world is supposed to be. To take Anscombe's famous example, a shopping list has a "world to list" direction of fit, but if a detective follows the shopper around, making notes about what he buys, these notes have a "list to world" direction of fit. Desires and intentions have a "world to mind" direction of fit, whereas beliefs and perceptual states have a "mind to world" direction of fit.

If, as I argue, pragmatic representations are both perceptual states and part of what constitutes the immediate mental antecedents of action, we are facing something like a dilemma. If they are perceptual states, they should have a "mind to world" direction of fit, like all perceptual states. But if they are the immediate mental antecedents of action, one could be tempted to think of their direction of fit as "world to mind"—like that of intentions.

One way of resolving this dilemma is to say that pragmatic representations have both directions of fit: they both describe how the world is and prescribe an action. In fact, those who explore the possibility of action-oriented or action-guiding representations that would mediate between sensory input and motor output, and that are to be distinguished from belief- or desire-like propositional attitudes, normally take this route (Millikan 1995, Pacherie 2000, 2011, Millikan 1995, Jeannerod and Jacob 2005, Clark 1995, Mandik 2005).

Millikan introduced the term "Pushmi-Pullyu representations" for representations that have both directions of fit: they represent the way the world is, and at the same time prescribe how the world is supposed to be. For example, the sentence "Dinner!" has such a double direction of fit: it both describes something (that dinner is ready) and prescribes an

action (that one should come to eat dinner). Millikan generalizes this account of Pushmi-Pullyu representations to mental representations. What Andy Clark calls "action-oriented representation," and what Jeannerod and Jacob call "visuo-motor representation," also have this double direction of fit (see especially Clark 1997, p. 50, Jacob and Jeannerod 2003, pp. 38, 204).[4]

I want to resist this move. I think it is problematic to talk about a perceptual state that prescribes what we should do (see more on this in Chapter 3). Further, I do not see any real reason for attributing "world to mind" direction of fit to pragmatic representations. We have seen that desires and intentions have a "world to mind" direction of fit. But pragmatic representations are neither desires nor intentions. Let me clarify.

Recall the distinction between the cognitive and the conative components of the immediate mental antecedent of action: the cognitive component represents the world, whereas the conative one moves us to act. As long as we make a distinction between these two components of the immediate mental antecedents of action, there is no reason why the representational component (what Brand calls the "cognitive" component) would need to have a "world to mind" direction of fit. The "conative" component moves us to act, and the representational component tells us how the world is in such a way that would help us to perform this movement. But are we entitled to make this distinction?

The advocates of the Pushmi-Pullyu approach could follow one of the following two strategies, in order to question this distinction between

---

[4] Pete Mandik's position is somewhat different from the others on this list. He claims that "even representations with only imperative content [that is, 'world to mind' direction of fit] are action-oriented representations" (Mandik 2005, Footnote 7). My view could be considered to be at the opposite end of the spectrum from Mandik's: while many proponents of the idea that action-oriented representations mediate between sensory input and motor output argue that these representations have both "mind to world" and "world to mind" direction of fit, Mandik claims that they can have only the latter. I claim they have only the former. Another philosopher who is difficult to characterize in this respect is Michael Wheeler (see Wheeler 2005, especially pp. 197–9). Sometimes he seems to take the "action-specific, egocentric, context-dependent" representations he is focusing on to describe how the world is: "how the world is *is itself encoded in terms of* possibilities for action" (Wheeler 2005, p. 197). But other times, he seems to take them to prescribe actions: "what is represented here is [. . .] knowledge of how to negotiate the environment, given a peculiar context of activity" (Wheeler 2005, p. 198). Maybe the most charitable interpretation is to treat these representations to have both directions of fit.

"cognitive" and "conative" components of the immediate mental antecedents of action. They could say that the representational component of the immediate mental antecedent of action is intrinsically motivating—hence, we don't need to posit two separate components. Or they could argue that the representational components of the immediate mental antecedent of action prescribes an action (although it does not intrinsically motivate us to perform it). As we shall see in Chapter 4 and Chapter 3, respectively, neither strategy is particularly promising.

I will say very little in this book about the "conative" component of the immediate mental antecedent of action (the component that moves us to act), and focus entirely on the "cognitive" one. In other words, this book is about the representational component of the immediate mental antecedent of action: about pragmatic representations. Pragmatic representations represent the world in a certain way: in a way that allows us, and guides us, to perform the action. The "conative" component of the immediate mental antecedent of actions is not easy to explain or understand—especially as the question of what "moves us to act" is so intertwined with the problem of free will. As a result, I leave the conative component aside. And my claim about the representational component of the immediate mental antecedent of action is that it is a perceptual representation. This is the claim I will defend in this chapter.

2.4 Pragmatic representations are perceptual states

What are pragmatic representations? My seemingly surprising answer is that they are in fact perceptual states.[5] I defined pragmatic representations as the representational component of mental states that make actions actions—that constitute the difference between actions and mere bodily movements, whatever they may be. I will now argue that if

---

[5] This proposal is not entirely new. Kent Bach (Bach 1978, p. 368), for example, claims, without argument, that one of the two components of the immediate mental antecedent of action (receptive representation) is perceptual, whereas the other component (effective representation) is "sensuous." See also my Nanay 2012a and 2011a, where I explicitly argued for a version of this claim, using different arguments.

we understand pragmatic representations this way, we have very good reasons to take them to be perceptual states.

Why should anyone be tempted by the view that pragmatic representations are, or even could be, perceptual states? One possible consideration is the following. If pragmatic representations "guide and monitor our ongoing bodily movement" (Brand 1984, p. 173), then this may suggest that they are perceptual states. Monitoring, if we can make sense of this concept in the present context, is a perceptual activity, and in order for the "guiding of ongoing bodily movement" to be efficient in real time it needs to be sensitive to immediate feedback about whether the bodily movement is succeeding.

Suppose you are trying to kill a mosquito in your hotel room. The mosquito is flying around, and you are following it around. Whatever representation guides your bodily movements has to be very sensitive to the momentary changes in the mosquito's whereabouts. If we consider pragmatic representation to be perceptual states, we have a straightforward account of how this happens: our bodily movements are sensitive to the whereabouts of the mosquito because the mental states that guide them are perceptual states—there is a direct feedback loop between what you do and what you see.

But note that this is only one way of describing this scenario. Another way of doing so would be to say that our perceptual states give rise to another, non-perceptual state, and the mental state that guides our bodily movement is this latter, non-perceptual state. Of course, this transition from the perceptual to the non-perceptual state would have to be a quick one, to allow for the real-time tracking of the mosquito.[6]

How can we decide between these two ways of describing the scenario? More generally, how can we decide between a perceptual and a non-perceptual account of pragmatic representations? I will argue in the next section that the latter account is inconsistent with an odd fact about perceptual learning.

---

[6] A further worry about these intuitive considerations is that the chasing of the mosquito could be thought to be a special kind of action, which is very different from other kinds of actions: for example, ballistic actions where there is no perceptual feedback at all. Hence, we should not generalize from the mosquito example to actions per se. I'll say more about ballistic actions in Section 2.6.

2.5 The argument

It is notoriously difficult to draw a clear distinction between perceptual states and non-perceptual mental states (see Siegel 2007, Kriegel 2007, Bayne 2009, Masrour 2011, Prinz 2006, Nanay 2010a, 2010d, 2011a, 2012a, Dennett 1996). Luckily, in order to argue against the view that pragmatic representations are non-perceptual states, we will not need a general criterion for keeping perceptual states and non-perceptual mental states apart. Suppose, for *reductio*, that pragmatic representations are non-perceptual states. We form them on the basis of perceptual states, but they are themselves not perceptual states. In short, each time we are performing a perceptually guided action we are in a perceptual state, P, and another non-perceptual mental state, NP, which is the immediate mental antecedent of our action and which we form partly on the basis of P. I say that we form NP *partly* on the basis of P because it is possible (indeed plausible) that NP is sensitive to some other non-perceptual states or our background knowledge.

Now consider the following short but impressive demonstration of perceptual learning.[7] Participants put on a pair of distorting goggles that shifts everything they see to the left: the scene in front of them appears more to the left than it is. Now they are supposed to throw a basketball into a basket. The first couple of attempts fail miserably: they throw the ball to the left of the basket. After a number of attempts, however, they are able to throw the ball accurately into the basket.

After having practiced this a few times with the goggles on, participants are asked to take off the goggles and try to perform the task without them. Now they go through the opposite perceptual learning process: when they first attempt to throw the ball towards the basket without the goggles, they miss it. After several attempts, they learn to throw it the way they did before putting on the goggles.

I would like to focus on the change in perception and action after taking off the goggles. The participants' pragmatic representations are clearly different at the beginning and the end of the learning process: at the beginning, the participants cannot throw the ball successfully into the

---

[7] This interactive demonstration can be found in a number of science exhibitions. I first saw it at the San Francisco Exploratorium. See also Held 1965 for the same phenomenon in an experimental context.

basket, but at the end they can. Their pragmatic representations change during this process, and it is this change that allows the participants to perform the action successfully at the end of the process. The mental state that guides their action at the end of the process does so much more efficiently than the one that guides their action at the beginning. Pragmatic representations represent the spatial location properties, the representation of which allows us to throw the ball. At the beginning of this process, they represent these properties incorrectly; at the end, they do so more or less correctly.

And here the proponents of the idea that pragmatic representations are non-perceptual states encounter a problem. Neither one's perceptual experience nor one's perceptual beliefs change in the course of this perceptual learning process. Background beliefs do not change either. The subjects experience the basket in front of them the entire time. Further, they have the same beliefs about the basket and its whereabouts all along. Still, their pragmatic representations do change.

The proponents of the idea that pragmatic representations are non-perceptual states are committed to the following picture. The pragmatic representation that guides this action is a non-perceptual state, call it NP. NP is based (partly) on one's perceptual state, P. Thus, the proponents of the idea that pragmatic representations are non-perceptual states need to explain this perceptual learning process in one of the following three ways, none of which would be particularly promising.

First, they can insist that NP is based entirely on P. In this case we get a flat contradiction, as while NP changes in this scenario, P, as we have seen, does not change. But if this is so, then NP cannot be based only on P—otherwise we reach a contradiction.

Second, they can point out, rightly, that we form NP only partially on the basis of P, and there are other non-perceptual states and background beliefs that also play a role in forming NP. But note that neither our background beliefs[8] nor our other (conscious) non-perceptual states

---

[8] One may wonder about whether background beliefs do remain the same. I certainly have acquired a belief after the first failed attempt; namely, the belief that I have just failed. Doesn't this count as a change in my background beliefs? Note that this change happens after the first failed attempt, which does not change the accuracy of my action: the second attempt normally fails just as miserably as the first. So although this is a change in my beliefs, this change cannot be construed as underlying the change in my pragmatic representation—which comes much later, after many more attempts.

change in this scenario: if neither P nor any other conscious mental states NP are supposed to be based on change, then it is difficult to explain why NP changes.

This takes us to the third option that is open to the proponents of the non-perceptual account of pragmatic representations. They could say that while our conscious perceptual states do not change, we have no reason to suppose that our non-conscious perceptual states do not change either. Maybe they do. In this case, besides the conscious perceptual state, P (that doesn't change), there is another perceptual state, call it P*, that is non-conscious and that, like NP, does change in the course of this process. The problem with this suggestion is that in this case, both P* and, presumably, NP would be unconscious representations of the subject. And we need to be careful about when we attribute unconscious mental states to agents. We have a clear and conclusive reason for attributing the pragmatic representation: the agent's behavior cannot be explained without attributing a representation of the basket's location that guides her action. But it is difficult to see what reason we have for attributing *two* unconscious representations to the agent besides salvaging the non-perceptual account.

The same objection applies in the case of yet another potential way of defending the non-perceptual account: namely, by arguing that while P does not change and NP changes, this does not pose any problem because there is another (non-conscious and non-perceptual) state that NP is partly based on, call it NP*, that does change. And this change escapes our attention because NP* is unconscious. The problem with this suggestion is the same as above: we have no non-ad-hoc reason to postulate yet another unconscious representation (perceptual or non-perceptual) in order to describe the agent's behavior and experience.

To sum up, none of the ways in which the non-perceptual account of pragmatic representations can explain the perceptual learning phenomenon proves to be satisfying—they either lead to straight contradiction or to the ad-hoc postulation of mental states. We are better off taking pragmatic representations to be perceptual states.

A possible final line of defense for the proponents of the non-perceptual account of pragmatic representation would be to say that the goggles/basketball scenario is unusual and we should not use it to draw general conclusions concerning the nature of the mental antecedents of actions. The problem with this response is that there

is an abundance of empirically documented cases where our pragmatic representations represent the world differently from the way our conscious perception does.

Here is an important set of examples. A number of optical illusions mislead our perceptual experience but not (or much less) our pragmatic representation. One such example is the three-dimensional Ebbinghaus illusion. The two-dimensional Ebbinghaus illusion is a simple optical illusion. A circle that is surrounded by smaller circles looks larger than a circle of the same size that is surrounded by larger circles. The three-dimensional Ebbinghaus illusion reproduces this illusion in space: a poker-chip surrounded by smaller poker-chips appears to be larger than a poker-chip of the same diameter surrounded by larger ones. The surprising finding is that although our perceptual experience is incorrect—we experience the first chip to be larger than the second one—if we are asked to pick up one of the chips, our grip-size is hardly influenced by the illusion (Aglioti et al. 1995, see also Milner and Goodale 1995, chapter 6 and Goodale and Milner 2004). Similar results can be reproduced in the case of other optical illusions, like the Müller-Lyer illusion (Goodale and Humphrey 1998, Gentilucci et al. 1996, Daprati and Gentilucci 1997, Bruno 2001), the "Kanizsa compression illusion" (Bruno and Bernardis 2002), the dot-in-frame illusion (Bridgeman et al. 1997), the Ponzo illusion (Jackson and Shaw 2000, Gonzalez et al. 2008) and the "hollow face illusion" (Króliczak et al. 2006).[9]

---

[9] I will focus on the three-dimensional Ebbinghaus illusion because of the simplicity of the results, but it needs to be noted that the experimental conditions of this experiment have been criticized recently. The main line of criticism is that the experimental designs of the grasping experiment and the perceptual judgment experiment are very different. When the subjects grasp the middle chip, there is only one middle chip surrounded by either smaller or larger chips. When they are judging the size of the middle chip, however, they are comparing two chips—one surrounded by smaller chips, the other by larger ones (Pavani et al. 1999, Franz 2001, 2003, Franz et al. 2000, 2003, see also Gillam 1998, Vishton 2004 and Vishton and Fabre 2003, but see Haffenden and Goodale 1998 and Haffenden et al. 2001 for a response). See Briscoe 2008 for a good philosophically sensitive overview of this question. I focus on the three-dimensional Ebbinghaus experiment in spite of these worries, but those who are moved by Franz et al.'s style considerations can substitute some other visual illusion; namely, the Müller-Lyer illusion, the Ponzo illusion, the hollow face illusion, or the Kanizsa compression illusion, where there is evidence that the illusion influences our perceptual judgments, but not our perceptually guided actions.

If we assume that pragmatic representations are non-perceptual states, then in this example the pragmatic representation NP that is responsible for the successful performance of the action would need to be formed (partly) on the basis of our perceptual state P. But P represents the size-property of the chips differently from the way NP does: the former represents it incorrectly, the latter correctly (or, to be more precise, less incorrectly). Thus, it is difficult to see how we can form NP on the basis of P—the assumption that pragmatic representations are non-perceptual states leads to very implausible consequences.

To sum up, the non-perceptual account of pragmatic representations does not seem particularly promising. It leads either to contradiction or to an ad-hoc postulation of mental states that we have no reason to make other than to salvage the non-perceptual account. But how can we account for the odd features of the perceptual learning scenario (and the three-dimensional Ebbinghaus case) if we accept my claim that pragmatic representations are perceptual states?

Let us take the goggles/basket scenario first. Our pragmatic representation attributes, perceptually, a certain location property to the basket, which enables and guides us to execute the action of throwing the ball into the basket, whereas our conscious perceptual experience attributes another location property to the basket. During the process of perceptual learning, the former representation changes, but the latter does not. How about the three-dimensional Ebbinghaus illusion? In this case, our pragmatic representation attributes, perceptually, a size-property to the chip, whereas our conscious perceptual experience attributes another size-property to it. Our conscious perceptual experience misrepresents the size of the chip, but the pragmatic representation represents the size of the chip (more or less) correctly.

Thus, we have two different perceptual states in the three-dimensional Ebbinghaus scenario: a conscious (incorrect) one, and a pragmatic perceptual representation. They are both representations: they both attribute properties to the same object. However, they attribute different properties. The conscious experience attributes the size-property we experience the chip as having. And the pragmatic representation attributes the size-property that guides our (successful) action. We have no reason to deny that we have both of these perceptual representations, and there is a lot of empirical evidence for the existence of two (more or less)

separate visual subsystems in humans that can represent the perceived object differently.[10]

2.6 Potential counterexamples

The claim that pragmatic representations are perceptual states is a surprising one, and I need to consider some potential *prima facie* counterexamples.

First, this proposal seems to suggest that one could not perform an action with one's eyes closed or in a dark room. Suppose that I wake up in my bedroom, which I know very well. It is pitch dark and I want to go to the bathroom. How could I perform this action? The representation that is supposed to guide this action is a perceptual state, but I do not perceive anything: it's too dark.

The first thing to notice is that I do perceive a lot of things with my tactile sense modalities, as well as by means of proprioception. I perceive whether I am lying on my back or on my side. I perceive where the bed ends.

Suppose now that I manage to get out of my bed. How can I perform the action of switching on the light or opening the door? The answer is that I have mental imagery of where the door is and where the light switch is. More generally, we can perform goal-directed actions on objects that we do not perceive as long as we have mental imagery of them. Put this book down in front of you and close your eyes. You can pick it up again without looking because you have a mental (presumably visual) imagery of where it is. Would this constitute a counterexample to the claim I have been defending?

Mental imagery is not strictly speaking a perceptual state. It is a quasi-perceptual state (on why and how mental imagery is quasi-perceptual, see Kosslyn 1994, Kosslyn et al. 2006, Page et al. 2011, Nanay 2010c).[11]

---

[10] Milner and Goodale 1995, Goodale and Milner 2004, Jacob and Jeannerod 2003, Jeannerod 1997. The dorsal visual subsystem is (normally) unconscious, and is responsible for the perceptual guidance of our actions. The ventral visual subsystem, in contrast, is (normally) conscious, and is responsible for categorization and identification. I do not want to rely on this distinction in my argument (but see Chapter 3 for a detailed analysis of how my account relates to the dorsal/ventral distinction).

[11] One may wonder about whether by taking imagery to be quasi-perceptual I am siding with Kosslyn (Kosslyn 1994, Kosslyn et al. 2006), and against Pylyshyn (see especially

So purists may want to read "quasi-perceptual" instead of "perceptual" in the argument I gave in the last section. I don't think much depends on this distinction. We have very good reasons to believe that mental imagery represents objects as having properties in the same way as perception does—as we will see in Chapter 5. Pragmatic representations that attribute properties to their objects by virtue of mental imagery, mental states I will call "pragmatic mental imagery," are important variants of pragmatic representation—so much so that I will devote an entire chapter (Chapter 5) to them.

Back to the counterexamples: how about scratching one's back without seeing it? What perceptual state guides this action? The simple answer, again, is that it is the tactile and proprioceptive perceptual state that represents the itch in a certain way and that guides the action (see Bach 1978 for a thorough analysis of the example of scratching an itch). Vision is not the only sense modality.

Third, is my argument valid for "ballistic actions"? Ballistic actions are actions that are performed without any perceptual feedback: throwing a ball as far as one can, jumping up as high as one can, and so on. What can we say about the immediate mental antecedents of ballistic actions? What properties do we need to attribute in order to perform these actions successfully? It has been noted that in the case of ballistic actions, although some of the properties of the objects we are performing actions with/on are irrelevant, many others are very relevant indeed (Kawato 1999, Seidler et al. 2004). When we are throwing a ball as far as we can, we do not need to represent where we are throwing it, but we do need to represent the weight of the ball, for example. When we are punching a bag as hard as we can, we do not need to represent the bag's weight or shape, but we need to represent its spatial location. In short, the mental antecedent of ballistic actions also guides the performance of these actions.

Pylyshyn 2002, 2003) in the grand "imagery debate" (Tye 1991). Note, however, that the imagery debate is about whether mental imagery is propositional or depictive. Even if one comes down on the propositional side, as long as one takes perception to be propositional (as, for example, Siegel 2010b does), this would be consistent with the claim that mental imagery is quasi-perceptual. In short, the "imagery debate" is orthogonal to the claim that mental imagery is quasi-perceptual, a (weak) version of which even Pylyshyn acknowledges (see especially Pylyshyn 2003).

Does this matter from the point of view of the applicability of my argument? No. The immediate mental antecedents of ballistic actions represent some properties of the object the action is directed at, and they do so perceptually. Now, it is not impossible that there are some ballistic actions the successful performance of which would not require the representation of any property of any object at all—although the standard examples, like ballistic jumping, are not among them (Hustert and Baldus 2010).[12] If there are, then the mental antecedent of these actions has only a conative, but no cognitive, component. But as my claim was that the cognitive component of the mental antecedents of actions is a perceptual state, these actions will be irrelevant from the point of view of my argument.

Finally, take the following two examples: blinking (without planning to do so, and without any external influence that would trigger the blinking reflex) and swallowing (again, without planning to do so, and without any external influence that would trigger the swallowing reflex). Are they actions? They are not intentional actions, but some may consider them to be actions nonetheless. They do seem to be different from reflexes (see Chapter 4). My bodily movement in the blinking case is not initiated by an external stimulus, as it would be in the case of the blinking reflex. Suppose, for the sake of argument, that they are *bona fide* actions.

Now, we need to remember the two different components of the mental antecedent of actions: the cognitive component that represents the world, and the conative one that moves us to act. In the case of blinking, it seems that the cognitive component is not present. We do not need to represent the world in any way in order to blink (or in order to swallow). Thus, in the case of these actions, the immediate mental antecedent of the action does not have a cognitive component. Hence, the question does not arise about whether this cognitive component is a perceptual state. Blinking and swallowing may count as a counterexample to the claim that the immediate mental antecedent of every action must include a cognitive component (a claim Brand and Bach seem to presuppose but I don't, see Chapter 4 below), but it does not qualify as a counterexample to the main claim of this chapter.

---

[12] Maybe the examples I consider in the next two paragraphs could be considered to be ballistic actions of this kind.

Final objection: how about complex actions? How about the action of finishing a book manuscript, for example? This is undoubtedly an action, but we certainly should not say that the immediate mental antecedent of this action is a perceptual state. My answer is that it is problematic even to talk about the immediate mental antecedents of complex actions like finishing a book manuscript. Finishing a book manuscript is an action that involves a lot of simpler actions—going to the office, turning on the computer, etc. All of these simpler actions have immediate mental antecedents, and, if I am right, they are all perceptual states, but the composite action of finishing a book manuscript does not have one single immediate mental antecedent. So the question about whether the immediate mental antecedent of complex actions like finishing a book manuscript is a perceptual representation does not arise.

2.7 Pragmatic representations everywhere

I argued that the mental antecedents of our actions are perceptual states. They attribute properties the representation of which is necessary for the performance of the action, and they do so perceptually: they guide our bodily movement. But if this is true then we may need to re-evaluate the way we think about perception. That is the task of the next chapter. We also need to adjust the belief–desire model of motivation. Some of our perceptual states—namely, pragmatic representations—play a more crucial role in bringing about our actions than has been acknowledged.

Sometimes what mediates between sensory input and motor output is not a set of beliefs and desires, as the classical belief–desire model suggests. Sometimes a pragmatic representation (given some appropriate conative mental state, which does not have to be a desire) is sufficient for bringing about an action. Further, even in those cases when beliefs and desires play an essential part in bringing about our actions, pragmatic representations still need to be involved.

Animals and small children are capable of performing goal-directed actions, such as running away from predators or chasing prey. But if they are, they must be able to have pragmatic representations. Hence, even organisms that may be incapable of entertaining complex thoughts and beliefs must be able to have pragmatic representations. These representations are basic, both phylogenetically and ontogenetically.

But if it is true that what mediates between sensory input and motor output for non-human animals and small children are pragmatic representations, then it is also true of the vast majority of the actions of adult human beings: most of the time, what mediates between our sensory input and motor output are pragmatic representations only—when we tie our shoelaces, empty the dishwasher, or just walk down the street without bumping into too many people.

The performance of *some* of our actions requires more complex and sophisticated mental processes, such as beliefs and desires, but these are only the tip of the iceberg. And, importantly, even the performance of these more sophisticated actions requires the fine tuning of pragmatic representations for them to be successful (Israel et al. 1993 and Perry 2001 make a similar point). I return to these questions in Chapter 4.

3

Perception

Seeing is the perception of action.

3.1 Are there perceptual representations at all?

I have been framing the discussion of what mediates between sensory input and motor output in terms of pragmatic representations. My main claim in the last chapter was that these pragmatic representations are perceptual representations. But are there perceptual representations at all? Twenty years ago, the vast majority of philosophers of perception would have agreed that there are, but this is no longer so. In fact, it seems to be one of the most widely discussed questions about perception these days (see Brogaard forthcoming a, for the current state of the debate).

I have been assuming here, and will be assuming in what follows, that perceptual states represent objects as having properties. But, and this is the rationale for the present section, my more general claims about the action-oriented nature of perception do not necessarily presuppose this representationalist framework.

Many of our mental states are representations: my belief that it is raining outside represents a putative state of affairs—that it is raining outside. If I am afraid of a tiger, this fear is also directed at, or is about, something: a tiger. In other words, many mental states refer to something, they are about something. Perceptual states are also about something: if I see a cat, it would be natural to say that my perceptual state is about this cat. It is tempting to conclude from this that perceptual states are also representations: they also have content. The cat (or some of the cat's properties), in this example, is part of the content of my perceptual state.

Some further *prima facie* reasons for being representationalist (see also Pautz 2010): it allows us to give simple explanations for illusions and hallucinations, as well as for perceptual justification. If we think of perceptual states as representations, then veridical perception amounts to having a correct perceptual representation, whereas illusion and hallucination amounts to perceptual misrepresentation (not everyone thinks this is an advantage, see Brewer 2006 for an analysis). And as perceptual states are at least sometimes capable of justifying beliefs, thinking of them as representations allows us to explain perceptual justification in very simple terms: as a relation between two different representations—a perceptual and a non-perceptual one (this is, of course, not the only way of explaining perceptual justification).

Thinking of perceptual states as representations is also a default assumption of (mainstream) perceptual psychology and vision science (see Burge 2005 and Nanay forthcoming c for summaries). But some philosophers (and some psychologists) are not convinced, and they conceive of perception in a non-representational manner.

Anti-representationalism is the view that there are no perceptual representations.[1] As it is a merely negative view, the question arises: what happens when we perceive, according to the anti-representationalist? There are a number of different suggestions, which fall into two broad categories: enactivism and relationalism.

According to relationalism, perceptual states are, in part, constituted by the actual perceived objects. Perception is a genuine relation between the perceiver and the perceived object—and not between the agent and some abstract entity called "perceptual content" (Travis 2004, Brewer 2006, 2011, Campbell 2002, Martin 2002b, 2004, 2006, forthcoming. See also Byrne and Logue 2008's and Burge 2005's criticisms, as well as Crane 2006). One reason why relationalism may seem appealing is that it captures the particularity of perception—the intuitively plausible assumption that the object of perception is always a particular token object—better than representationalism (see Soteriou 2000 for a summary).

---

[1] Some anti-representationalists make a more modest claim: that *veridical* perceptual states are not representations, while allowing for non-veridical perceptual representations (see Pautz 2010 for discussion).

Suppose that I am looking at a pillow. What happens if someone, unbeknownst to me, replaces this pillow with another, indistinguishable, pillow? Most representationalists will say that my perceptual state is still the same, as this replacement does not make a difference to the content of my perceptual state (note that not all versions of representationalism are committed to this claim: those, for example, that conceive of content as Russellian, gappy, or "singular when filled"—see, for example, Tye 2007, Dokic 1998, Schellenberg 2010—are not). But I am looking at two entirely different token objects before and after the swap. The relationalist thus insists that I have two completely different perceptual states.

Like relationalists, enactivists also deny that there are perceptual representations, but they give a different (but not incompatible, see Noë 2004, Hellie forthcoming) positive account of perception. According to one version of enactivism, perception is an active and dynamic process between the agent and the environment, and this dynamic interaction doesn't have to be (or maybe even couldn't be) mediated by static entities like representations (Chemero 2009, Port and van Gelder 1995). Another version of enactivism emphasizes that when we see a scene, the whole scene in all its details is not coded in our perceptual system. Only small portions of it are: the ones we are attending to. The details of the rest of the scene are not coded at all, but they are available to us all along—we have "immediate perceptual access" to them without representing them (O'Regan 1992, Noë 2004, especially pp. 22–4).

One may wonder whether these enactivist claims about the nature of perception give us reason to abandon the idea of perceptual representations per se, or maybe only to conclude that they are not static or not detailed. In short, some of the enactivist arguments may give us good reason to prefer certain kinds of perceptual representations over others within the representationalist framework. But that is clearly not the aim of most enactivists, who want to reject the whole idea of perceptual representations.

The debate about perceptual representation is a subtle one, and both sides should be taken seriously. I elsewhere offer four possible ways of resolving this debate by (a) capturing some anti-representationalist intuitions within the representationalist framework (Nanay 2012d), by (b) discrediting anti-representationalism on empirical grounds (Nanay forthcoming c), by (c) exploring the possibility that the two camps are talking about different phenomena (Nanay forthcoming c), and by

(d) finding a framework where the two views can coexist as different explanations for different explanatory projects (Nanay forthcoming f). But what matters for us here is that regardless of which side of the debate we take, the action-guiding nature of some of our perceptual states still needs to be emphasized and explained.

The main claim of this book—that there are special mental states that mediate between sensory input and motor output—is an important one in either the representational or the anti-representationalist framework. If the reader is drawn to representationalism, she can just read ahead: I will continue to talk about perceptual representations and I will talk about pragmatic representations as perceptual representations. But if the reader is drawn to some version of anti-representationalism, this is not a good reason to stop reading this book.

Many, maybe even most (although probably not all), of the claims I am making about pragmatic representations can be rephrased in an anti-representationalist framework. I will spend much of this chapter specifying what properties are represented in perception: what properties perceptual representations attribute to objects. The anti-representationalist can read this discussion as being about what properties' perceptual states "present" or "are sensitive to" or "track." Those who reject the claim that perceptual states are representations nonetheless still think of perceptual states as being in a perceptual relation with objects and their properties. The question that I frame as "What properties are represented in perception?" could be framed for these anti-representationalists as "What properties of perceived objects are we perceptually related to?"

3.2 Perceptually attributed properties

Again, I will assume for the sake of simplicity in what follows that there are perceptual representations. The question now is this: What are they? There are two very different approaches to characterizing perceptual representations (these are not two theories, but rather two kinds of general approaches).

The first one is to start out with non-perceptual representations, typically beliefs, and see how what we know about representations of this kind can be modified in order to apply to the perceptual case. Some

think that there is no need for any modification: perceptual content is exactly the same as belief content. But most philosophers who think of perceptual content this way allow for some differences—while nonetheless maintaining that we should use propositional content as a model for understanding perceptual content.

Much of these proposed modifications aim to address the problem of the particularity of perception that I mentioned above. The general idea is that unlike the content of beliefs, perceptual content somehow depends constitutively on the token perceived object. However, these "Russellian," "gappy," "singular," "object-involving" or "singular-when-filled" conceptions of perceptual content are nonetheless conceptions of propositional content—as David Chalmers says, these accounts are thinking about perceptual content as a "structured complex" (Chalmers 2006, p. 54. Thompson 2009 describes them even more aptly as "structured propositions").

The second approach to characterizing perceptual representations is to resist the temptation to start out with belief content, and instead use a more basic way of thinking about content in general that can subsume both belief content and perceptual content. We have no reason to believe that all mental representations are linguistically or propositionally structured (see Crane 2009, Burge 2010, but see also Siegel 2010a, 2010b). Some (but not all) mental states have content. Some of these (but not all of them) have propositional content. But perceptual states don't.

I vastly prefer the second approach, although it needs to be acknowledged that the differences between the two very much depend on one's conception of propositional content. As we have seen in Chapter 1, not everyone thinks that propositional content is syntactically or linguistically structured. Some think of propositions as sets of possible worlds (Stalnaker 1976). If we think of propositions in this weak sense (and if the talk of content as sets of possible worlds is unproblematic, which is not an obvious assumption), then I have no problem taking perceptual content to be propositional. But if propositions are taken to be syntactically or linguistically structured, then I would vastly prefer not to think of perceptual content this way. Without attempting to settle this terminological debate about propositions and propositional contents, I want to use a general enough way of thinking about perceptual content that would allow us to characterize the content of pragmatic representations.

What would then be a general enough way of thinking about mental representations in a not necessarily propositional (that is, syntactically structured propositional) manner? A reasonable suggestion is to think of them as attributing properties to entities. And if we think of mental representations in general as attributing properties to entities, then we should think of perceptual representations as perceptually attributing properties to the perceived scene.

Two major questions remain:

(a) What are these properties?
(b) What is the "perceived scene"?

The hope is that by focusing on pragmatic representations, we can make progress in understanding both of these aspects of perceptual representations. I address perceptually attributed properties in this Section and turn to the question of the "perceived scene" in Section 3.4.

What are these properties that are perceptually attributed when we perceive? One way of understanding this question is to interpret it as asking whether these properties are tropes or universals (Nanay 2012d, Campbell 1990, Mulligan 1999, Mulligan et al. 1984). Another way in which the nature of the perceptually attributed properties needs to be specified, and one I will return to briefly in Chapter 5, is whether they are determinates or determinables (or maybe super-determinates). The question I want to focus on here is the following: What is the range of properties that are perceptually attributed?

Beliefs can represent their objects as having any property. Perceptual states, in contrast, represent their objects as having a limited set of properties. Some plausible candidates include having a certain shape, size, color, and spatial location. The list may be extended, but it will not encompass all properties. The property of having been made in 2008 in Malaysia is unlikely to be represented perceptually—it is a property that is likely to be attributed by a non-perceptual state.

The question is then: Which properties are represented in perception and which ones are not? A couple of quick examples: it has been argued that we perceive objects as trees and tables (Siegel 2006a), as being causally efficacious (Siegel 2005, 2006b, 2009), as edible, climbable, or Q-able in general (Nanay 2011a, 2012a), as agents (Scholl and Tremoullet 2000), as having some kind of normative character or value (Kelly 2010, Matthen 2010a), as having dispositional properties (Nanay 2011b),

as having moral value (Kriegel 2007), and as affording certain actions (for very different versions of this claim, see Gibson 1966, 1979, Bach 1978, especially p. 368, Jeannerod 1988, 1994, especially Section 5, 1997, Jacob and Jeannerod 2003, especially pp. 202–4, Humphreys and Riddoch 2001, Riddoch et al. 1998, especially p. 678).

Depending on our view on what range of properties we attribute perceptually, we end up with very different views about perceptual content and, as a result, about perception in general. What is important for the purposes of this book is that if we take pragmatic representations to be perceptual states, then we get a somewhat surprising answer to the question about the range of perceptually attributed properties.

3.2.1 Action-properties

We have seen that pragmatic representations attribute properties, the representation of which is necessary for the performance of an action. And I have argued that pragmatic representations are perceptual states. If we put these two claims together, what we get is that we perceptually represent those properties the representation of which is necessary for the performance of an action. I call these properties, the representation of which is necessary for the performance of the subject's action, "action-properties." The claim then is that pragmatic representations attribute action-properties perceptually. But what are these action-properties? What properties are such that the representation of them is necessary for the performance of an action?

Let us return to the example of my action of picking up a cup of coffee. What are the properties of this cup that need to be represented in order for me to be able to perform this action? I need to represent its spatial location, its size, and its weight. Moreover, I need to represent these properties in an action-relevant way: in a way that would help me to pick it up. It would be of no use for the purposes of performing this action if I represented the cup's weight in pounds or its size in inches—in any case, such a representation would not be necessary for the performance of the action. The representation of the cup's weight in such a way that would allow me to exert the appropriate force is, in contrast, necessary.

Further, action-properties are relational properties: they depend both on the object and on my strength, hand size, and spatial position, in this example. To focus on the size-property only, the action-property that is necessary for performing the action of lifting the cup is a relational

property, the *relata* of which are the size of the cup and the grip size I would be approaching the cup with, which, in turn, depends on the size of my hand.

It is an important part of the definition of action-properties that they are properties of the object in front of us—not properties of our hand or our movement. As we have seen, action-properties are relational properties with two *relata*: the object's property and some property of myself or my own behavior (in this case, the size of the cup and my own grip size). But they are attributed to the object—the cup, not my hand. Attributing a property to the cup—of having a certain size related to my grip size—is different from attributing a property to my hand of having a certain grip size related to the cup's size, in the same way as attributing the property of being left of the pine tree to the cedar is different from attributing the property of being right of the cedar to the pine tree. The property of having a certain grip size related to the size of the cup, a property attributed to my hand, is not an action-property, as I defined action-properties. But it is important to emphasize this difference between properties of this kind and action-properties.

According to Henri Poincaré, "to localize an object simply means to represent to oneself the movements that would be necessary to reach it" (Poincaré 1905/1958, p. 47). Poincaré's proposal is about an important constituent of all perceptual states, not just the ones that are directly involved in guiding actions: in order to even localize an object, one needs "to represent to oneself the movements that would be necessary to reach it." But even if we ignore this difference in scope, and limit his claim to perceptual states that are directly involved in guiding actions, his claim is very different from mine. Suppose that we weaken his claim in the following way: in order to perform the action of reaching for something, we need to "represent to oneself the movements that would be necessary to reach it." This is still very different from my claim. Representing to oneself the movements that would be necessary to grasp the coffee cup is, as we have seen, different from representing the cup in such a way that would allow me to grasp it.

A similar proposal is supplied by Kent Bach, who argues that one of the constituents of the immediate mental antecedents of action (which he calls "effective representation") is the representation "of immediately subsequent behavior, whether or not it occurs" (Bach 1978, p. 367). Again, the representation of immediately subsequent behavior

is different from the representation of the object as having properties the representation of which is necessary for performing the action.

I emphasized the differences between the representations Poincaré and Bach talk about and pragmatic representations, but maybe their accounts and mine can be brought closer to one another. If, as Poincaré and Bach suggest, the immediate mental antecedents of action represent the agent's movement or body, they normally (maybe not necessarily) also represent the properties of the object on which the action would be performed. The representation of "the movements that would be necessary to reach" the object normally presupposes a representation of the object itself—and all of its properties that are relevant to the performance of this movement.[2] If I did not represent the spatial location and size of the object, it would be difficult to represent "the movements that would be necessary to reach it." So if we accept the suggestion of Poincaré and Bach, it seems to follow that we need to represent objects as having action-properties. But not *vice versa*. We can represent the cup as having a size-property that helps me to reach for it with the appropriate grip size without representing the movement that I would be approaching it with. Hence, those who are tempted by the approach of Poincaré and Bach should be able to accept everything I will say in what follows.[3]

Another approach to pragmatic representation that I need to separate my account from is the view that the "representational underpinnings of action" (Butterfill and Sinigaglia forthcoming, Section 1) are the representations of the outcome of the action (Millikan 2004, Chapter 16, Butterfill and Sinigaglia forthcoming). The view is that when I pick up a cup, what needs to be represented is the outcome of the action—the cup in my hand. Note that this view is very different from mine: all my view requires is that the cup's shape, size, weight, and spatial location

---

[2] It is not always clear what Poincaré means by the "localization of an object." One possible interpretation is that it means what I mean by the attribution of spatial location properties. If this is so, then his account and mine do not seem to be reconcilable: he explains the attribution of action-properties by means of the representation of one's own movement, and I explain the representation of one's own movement (if there is such representation) by means of the attribution of action-properties. One possible argument on my side (that I won't explore here) is that if Poincaré takes this route, he has no way of explaining what makes a goal-directed action goal-directed.

[3] There may be some empirical reasons that militate against the Poincaré/Bach view (see Butterfill and Sinigaglia forthcoming, especially Section 2). These arguments do not apply in the case of my account.

properties are represented.[4] I do not want to exclude the possibility that the outcome of the action *can* be represented (and there are some empirical reasons to think that they, at least sometimes, are in fact represented, see Bonini et al. 2010, Ambrosini et al. 2011), but it is not necessary to represent it in order for the action to be performed. The representation of the object's action-properties, in contrast, is by definition necessary for the performance of the action. Further, what may sound like a harmlessly simple representation of goals is in fact not so harmless, as it involves the representation of a counterfactual situation— I see the window closed and I represent it opened (by me). So if we took this route, it would make it somewhat difficult to talk about the actions of animals that are incapable of representing counterfactual situations. I consider this to be a major drawback of the goal-representation approach.

A couple of further clarifications are in order about action-properties and the perceptual representation thereof. First, we may perceive an object as having a certain action-property, but the object may fail to have this property. Our perceptual state may misrepresent action-properties. Second, we may perceive the same object as having different action-properties. I can, for instance, perceive a newspaper as having various action-properties in different contexts: I can perceive it as having an action-property that is relevant to my action of killing a fly, an action-property that is relevant to my action of reading about the election results, etc. The action-properties we attribute to the newspaper—that is, the content of pragmatic representations—can be sensitive to the action we are inclined to perform.[5]

Third, perceiving action-properties is not sufficient for the performance of actions. Often, the agent perceives an object as having an action-property, but the action itself is not performed. As we have seen, pragmatic representations form the representational component of the immediate mental antecedents of action. The immediate mental

---

[4] Interestingly, Marc Jeannerod seems to be oscillating between these two views: between "representation of goals for actions" and "visuomotor representations" (Jeannerod 1994, Section 5, Jeannerod 1997, Jacob and Jeannerod 2003, pp. 202–4).

[5] This is not to say that it is always sensitive to the action we are inclined to perform. My claim does not exclude the possibility that in some cases, we can perceptually represent an action-property even if we did not previously intend to perform an action (see, for example, Ellis and Tucker 2000).

antecedents of action also have a conative component, and if this conative component is missing, the action is not performed. I will say more about this in the next chapter.

Finally, the perceptual representation of action-properties, just like any perceptual state, can be conscious or unconscious. In other words, we may or may not be aware of the action-properties that our pragmatic representations represent objects as having. We can perform actions without consciously perceiving anything as having any action-properties if, for example, we act without attending to what we are doing. But in order for us to perform this action, we must perceptually represent action-properties: we are just not doing so consciously.

3.2.2 Thick action-properties

According to the picture I outline in this chapter, action-properties are typically attributed unconsciously: by typically unconscious pragmatic representations. But sometimes we are aware of those properties of an object that are relevant to our action. This subsection is about the properties that we experience the objects we act on as having.

The question then is: what are the properties that we experience objects as having when we have a pragmatic representation? Are the properties we experience, and the properties our pragmatic representations perceptually represent, the same? There is no reason why they should be the same. While our pragmatic representation attributes action-properties to an object, we may experience the object as having much richer properties. Are these richer properties attributed by the pragmatic representation itself? I want to remain neutral about this. If they were attributed by pragmatic representations, then, given that pragmatic representations are perceptual states, these attributed properties would be part of our perceptual phenomenology.

But perceptual phenomenology is notoriously difficult to differentiate from non-perceptual phenomenology (Masrour 2011, Siegel 2007, Kriegel 2007, Bayne 2009, Nanay 2011a, 2012e). Here is an example: at a dinner party, I'm eating a piece of meat that I take to be chicken, when my host tells me that it is in fact a piece of rat meat (or pigeon, etc.—use your favorite disgusting animal). The phenomenal character of my experience changes, but is this change a change in my perceptual (sensory) phenomenology or in my non-perceptual phenomenology? The answer is not at all clear. Further, if my introspection tells me that it is

a change in my perceptual phenomenology and your introspection tells you that it is not, how can we settle this disagreement?

Because of these difficulties surrounding perceptual phenomenology, I will not assume here that all the properties we experience objects as having when we have a pragmatic representation are perceptually represented (but I did give an empirical argument that is supposed to show that thick action-properties are part of our perceptual phenomenology in Nanay 2012e). But this means that I am not in a position to assume that they are represented by the pragmatic representation itself either. It is possible that these properties are represented by another perceptual state—in the way that the size-property we experience the chips as having in the three-dimensional Ebbinghaus illusion is represented by another perceptual state—not by the pragmatic representation. It is also possible that on the basis of your pragmatic representation, you form a non-perceptual state that attributes these properties you experience the object as having. This account may or may not presuppose some kind of non-sensory or "cognitive" phenomenology, and I do not want to endorse any such account. As I want to remain as neutral towards philosophical debates about phenomenology as possible, all I will be assuming in what follows is that when your pragmatic representation attributes action-properties to an object, you may experience this object as having a range of properties. Just what mental state attributes these properties is a question I want to bypass.

Can we be aware of the action-properties themselves? Yes, we can: if we attend to the weight property of the cup while lifting it, in order to exert the appropriate force, we consciously represent this action-property. As we have seen, this representation does not have to be conscious, but it can be—especially when we are performing some unusual action that we are not used to performing.

But the properties we experience an object as having when we have a pragmatic representation can be very different from, and, in fact, much richer than, the properties our pragmatic representation (unconsciously) represents this object as having. Here is an example. You have to wait in a very cold waiting room at the doctor's office. Suddenly, you spot the heater unit that is switched off. You walk to the unit and turn the knob.

The pragmatic representation that allows you to do so represents the size, shape, and spatial location properties of the knob. It represents the knob as having a certain shape that allows you to approach it with the

appropriate grip size, a certain spatial location that allows you to reach in the appropriate direction, and so on. These are the action-properties your pragmatic representation represents the knob as having. It does not represent anything about what effect the turning of the knob would produce.

But you can, and often do, experience the knob as having much richer properties: for example, the property of being a facilitator of turning on the heat. These properties are not action-properties. Remember, action-properties were properties the representation of which are necessary for the performance of the action. The property of being a facilitator of turning on the heat is not a property of this kind.

We may experience the knob as having all kinds of other properties: as being red, for example. I am interested in a subset of the properties we experience the knob as having that, like action-properties, cannot be fully characterized without reference to the agent's action. I call properties of this kind "thick action-properties." Thick action-properties, as we have seen, can be very different from, and much richer than, action-properties themselves.

Again, thick action-properties are the properties one consciously attributes to objects that cannot be fully characterized without reference to one's action. And one's action can characterize these consciously attributed properties in a variety of ways: one can experience objects as affording actions, as being an obstacle to one's action, as being a facilitator of one's action, as something one can perform the action with, as something one can perform the action on, as something that one should perform the action with, and so on. We can be aware of thick action-properties even if we are not performing any action. If, for example, I see a tiger running towards me, I will presumably experience it as having thick action-properties, but I may not perform any action at all.

I said in Chapter 2 that my discussion of pragmatic representation will focus on basic actions (actions that are not performed by means of the performance of another action), rather than on non-basic actions (actions that are performed by means of the performance of another action). I also said that the pragmatic representation that guides my basic action is exactly the same as the pragmatic representation that guides my non-basic action (i.e., the action that is performed by means of the performance of this basic action). But the associated thick action-properties can be very different (see Witt et al. 2005, Nanay 2006a).

An example: the pragmatic representation that makes it possible for me to press enter on my laptop is the same pragmatic representation that formats the hard drive as a result. They both represent action-properties that allow me to do so. But I will experience the enter key in these two situations very differently: as having very different thick action-properties. In the case of the performance of the basic action, I presumably experience it, if I experience it at all, as a key I am about to press (of a certain shape, at a certain spatial location). In the case of the non-basic action, I presumably experience it as the facilitator of some major, and potentially disastrous, procedure: as the means to erase all the data on my hard drive.

Thick action-properties can be emotionally charged. When I am about to press enter in the previous example, I may experience it as scary, as risky, as frightening, etc. If I suddenly notice a slug climbing up my bare left foot, I form a pragmatic representation that guides my action of getting rid of it. I also experience the slug as having thick action-properties, and these thick action-properties are very likely to involve disgust. But, again, this does not imply that action-properties themselves are emotionally charged.

When we are performing an action, we must have a pragmatic representation. And if we are being aware of performing an action, we experience the objects involved as having thick action-properties. If we attribute action-properties, and we are aware of doing so, then we also attribute thick action-properties. Can we be aware of thick action-properties without having a pragmatic representation? Arguably, we can. Suppose that you are trying to diffuse a bomb that will explode in 60 seconds. You have no idea how to do that, so you can't form a pragmatic representation that would specify the properties you need to represent in order to perform an action. But you do very much experience the bomb as having all kinds of thick action-properties. Your experience will be very much about properties that cannot be fully characterized without reference to some of your actions.

Thick action-properties can be thought to be the experiential counterpart of action-properties, but it is important to note that action-properties proper have to be specific enough so as to help us guide the actual performance of a specific action. There is no such requirement on thick action-properties: they can be very unspecific—as in the case of

diffusing the bomb. Thick action-properties specify actions on a much more general level than action-properties proper.

Properties like "edible" and "climbable" are thick action-properties—not action-properties proper (I myself have not been careful enough in the past to draw this distinction). So is the property of "affording a certain action." The properties that "causal indexical" predicates express—like "too heavy," "within reach," "is a weight I can easily lift" (Campbell 1993, especially pp. 82–8, see also Campbell 1994)—are also thick action-properties. What Siegel calls "experienced mandates" (Siegel forthcoming) are also forms of thick action-properties. Thick action-properties can be, and often are, dispositional properties.[6] They can also be, as we have seen, normative and emotionally charged. But action-properties proper are neither normative nor emotionally charged. Much of this book is about the attribution of action-properties—about pragmatic representations—but in Chapters 5 and 6 I say more about various forms of attributing thick action-properties.

3.2.3 Action-properties are not normative

Thick action-properties can be, and often are, normative: we can really experience an object as inviting eating, or as pulling us to perform some action or another with it (see Cussins 2003, Kelly 2010, Dreyfus 2005 and Siegel forthcoming for a critical analysis). And the emphasis on this experiential aspect of performing actions has an impressive history, going back to at least the Gestalt psychologists, who talked about "demand characters"—aspects of the world that put demands on us. As Kurt Koffka says, "to primitive man each thing says what it is and what he ought to do with it [. . .] a fruit says 'Eat me'; water says 'Drink me'; thunder says 'Fear me'" (Koffka 1935, p. 7). Gibson's concept of affordance (Gibson 1979) arguably picks up on this idea (although Gibson himself seems to oscillate between two interpretations of "affordances"—the normative "demand character" one, and a weaker non-normative "action-possibility" one).

---

[6] Can action-properties proper be dispositional properties? If they can, this would mean that dispositional properties can be perceptually represented. I am not committed to this conclusion here, but as I pointed out in Nanay 2011b, the recent mistrust about the possibility of the perceptual representation of dispositional properties is unfounded.

"Demand character" is not an action-property in my sense. The "demand character" of objects, however, would qualify as a thick action-property: we do experience objects as having this property, and characterizing this property involves a necessary reference to one's action. But, again, it is not an action-property proper. We should not confuse thick action-properties with action-properties proper. What thick action-properties we experience objects as having when we have a pragmatic representation says very little about what properties these (typically unconscious) pragmatic representations represent objects as having. Just because "demand character" is a thick action-property, we have no reason to take it to be an action-property proper.

More generally, I am very suspicious of the idea that our perceptual states attribute these "demand characters" to objects. Part of the problem is that if this were the case, then our perceptual states (as well as the perceptual states of simple animals) would attribute normative properties—a claim many philosophers, and especially vision scientists, would surely resist. But even those who are tempted to accept it (see, for example, Kelly 2010 and Matthen 2010a for recent defenses of somewhat similar ideas) will find it difficult to construe the properties the representation of which is necessary for the performance of an action as normative, or as "demand characters." These properties are the properties of objects: the location of the cup, its size, its weight. Although these properties are represented in an action-relevant way, we have no reason to suppose that the cup is represented as something we *should* lift.

Remember the definition of action-properties: properties the representation of which is necessary for the performance of the agent's action. The coffee cup's action-properties are the properties that I have to represent in order to be able to perform an action with it—in order to lift it, for example. Representing the coffee cup's "demand character" is not necessary for the performance of my action. It is enough if I represent its size, weight, and spatial location in such a way that would help me to lift it. I don't have to represent what I should do with it. In other words, representing the coffee cup as having normative properties is superfluous. And it is also suspicious from an evolutionary point of view: why would evolution have bothered to enable us to represent complex normative properties, when all that is required for performing an action is to represent those properties of the object that allow one to perform the action?

These problems with the proposal that pragmatic representations attribute normative properties to objects will become especially important if we return to the question of the direction of fit of these representations. As we have seen, several philosophers argued that the immediate mental antecedents of action do not "describe" how the world is, but rather "prescribe" how it should be—or, that they do both: they describe and prescribe. They have the direction of fit typical of desires and intentions, not of (or in addition to that of) beliefs and perceptual states.

I see no reason why we should accept that pragmatic representations "prescribe" what we *should* do with the object in front of us. My pragmatic representation attributes a size-property to the cup in front of me. Without representing this property, I would not be able to perform the action of lifting the cup. The function of pragmatic representations is to specify the size, location, and weight of the cup in an action-relevant way: in a way that can be used for the purposes of performing the action of lifting it. Their function is not to tell us what we should do with the cup, but to specify some of its properties for the purposes of acting on them.

3.3 Sensory individuals

We perceive things around us as having various properties: we perceptually represent certain entities as having certain properties. Let us return to the two questions that need to be answered about this picture: what are these entities and what are these properties? More precisely:

(a) What does our perceptual system attribute properties to?
(b) What kinds of properties does our perceptual system attribute to these entities?

I addressed (b) in the last section. It is now time to turn to (a). We have seen that not all properties that we represent objects as having are perceptually represented. If I am looking at an apple, I presumably perceive it as red and as having a certain size, shape, and spatial location, but what other property-types do I perceive it as having? I may represent the apple as having the property of being a Granny Smith apple picked by Mr. Jones in Oregon in 2014, but these are not properties that are likely to be perceptually represented. We need to have a distinction between

properties of the perceived object that are perceptually represented, and properties that are non-perceptually represented. I also argued that action-properties are perceptually represented.

Note that the same structural problem arises when we try to answer (a): what kind of entities are the ones these properties are perceptually attributed to? We may attribute properties non-perceptually to everyday objects, like the apple or the cedar tree in front of my window, but the question is what we perceptually attribute properties to. As the range of properties attributed perceptually and non-perceptually are often not the same, the kind of entities these properties are attributed to may not be the same either.

Question (a) needs to be distinguished from the old debate about the object of perception—that is, about what physical objects our perceptual system tracks (Clarke 1965, Strawson 1979). Question (a) is about what our perceptual system represents, and not about the metaphysical status of entities out there. If we accept the consensus view that the objects of perception are physical objects, this does not mean that our perceptual system represents physical objects as having various properties. To avoid confusion, I call the entities our perceptual system attributes properties to "sensory individuals" (I borrow the term from Cohen 2004). We can rephrase (a) as "What are sensory individuals?"

3.3.1 The sensory individuals of vision

Why should we posit sensory individuals at all? The widely cited reason is that properties need to be bound to particulars in perception. To use the classic example (Jackson 1977), if I see a red triangle and a blue square, the property of being red and of being a triangle are bound to one particular, whereas the property of blue and of being a square are bound to the other. If they were not bound, then this perceptual state would be indistinguishable from seeing a red square and a blue triangle. This consideration only specifies that the perceptually attributed properties need to be bound to a particular, but there is more than one candidate for what this particular may be.

A sensible suggestion would be to say that we perceive ordinary objects. When I see an apple as red, I perceptually attribute the property of being red to the apple. As David Armstrong says, "In perception, properties and relations are attributed to objects" (Armstrong 2004, p. 20, see also Shoemaker 1990, p. 97, Brewer 2007, p. 88, to mention

just a few examples). The concept of "ordinary object" needs to be clarified, as the mainstream conception of sensory individuals should not rule out shadows and rainbows, which are not physical objects. Here is Mohan Matthen's definition: the perceived object is a "spatiotemporally confined and continuous entity that can move and take its features with it" (Matthen 2005, p. 281, see also Pylyshyn 2007, Cohen 2004, Matthen 2004, 2010b for similar views).

The other main (contemporary) candidate for sensory individuals would be spatiotemporal regions. The main champion of the alternative, minority, view is Austin Clark, who characterizes sensory individuals (which he calls "phenomenal individuals") as "regions or volumes at which qualities seem to be located" (Clark 2000, p. 61). In other words, according to Clark, our perceptual system attributes properties to places, not ordinary objects. As he says, "the sensation of a red triangle [. . .] picks out places and attributes features to them" (Clark 2000, p. 147, see also Clark 2004 for clarifications).

Thus, we have two competing views for answering (a):

- The "ordinary object" view: sensory individuals are "spatiotemporally confined and continuous entit[ies] that can move and take [their] features with [them]" (Matthen 2005, p. 281)
- The "spatiotemporal region" view: sensory individuals are "regions or volumes at which qualities seem to be located" (Clark 2000, p. 61)

The "spatiotemporal region" view is widely considered to be wrong. Here are four arguments against it:

  (i) The "spatiotemporal region" view is too revisionary: we take ourselves to be perceiving ordinary objects: tables, water bottles, trees. If the "spatiotemporal region" view is correct, we're almost always wrong about what we perceive (Cohen 2004, p. 476, see also Matthen forthcoming).

 (ii) If the "spatiotemporal region" view is correct, then it is difficult to account for perceptual justification: perception is about spatiotemporal regions, whereas beliefs that are perceptually justified are about ordinary objects. As Susanna Siegel summarizes, "if audition told us that it was a place, rather than something at that place, that was cheeping, we would have all sorts of errors to correct in the move from audition to thought" (Siegel 2002, p. 137, see also Matthen forthcoming).

(iii) The "spatiotemporal region" view has trouble explaining how we can perceptually track an object over a period of time. Here is Siegel again: "What happens in sensory phenomenology when a subject sees a basketball make its way from the player's hands to the basket? The information that it's one and the same basketball traversing a single path is not given by sentience if sentience is limited to feature-placing. [According to the 'spatiotemporal region' view], the information that it's one and the same basketball traversing a single path has to be given non-sensorily. The subject's visual experience stops short" (Siegel 2002, p. 137, see also Matthen 2004, 2005, Cohen 2004, Pylyshyn 2007, p. x). Note that this is especially disturbing in the face of evidence that one-month-old babies can also track moving objects through space (see Johnson et al. 1991 for a summary) and, according to Siegel's argument, this would imply that they do so non-perceptually.

(iv) Finally, and perhaps most convincingly, there is empirical evidence against the "spatiotemporal region" view: it turns out that we can track the changes of two different entities at the very same spatiotemporal region (as Blaser et al. 2000 demonstrated with Gabor patches)—we seem to be tracking the changes of two sensory individuals, but they are at exactly the same spatiotemporal region. Hence, sensory individuals cannot just be spatiotemporal regions.

Note that almost all of these arguments are about vision. And they may indeed give us conclusive reason to accept the "ordinary object" view for the visual sense modality. But it is far from clear that this conclusion generalizes for the other sense modalities. A very commendable trend in contemporary philosophy of perception is the shift of emphasis from vision to other sense modalities: it seems clear that very often it is a bad idea to generalize from vision to the other sense modalities (O'Callaghan 2009, 2011, Lycan 2000, Batty 2010, Clark 2011). And a good example for such discrepancies between different sense modalities concerns sensory individuals.

3.3.2 The sensory individuals of audition

The debate about the sensory individuals of audition is very different from the "ordinary object" versus "spatiotemporal region" debate concerning

the visual sense modality. The main candidate for the sensory individuals of audition are sounds. Here is a typical statement from Casey O'Callaghan:

What do we hear? Sounds are, in the first instance, what we hear. They are the immediate objects of auditory experience in the following sense: whatever else we might hear, such as ordinary objects (bells, trumpets) and events (collisions, typing), we hear it in virtue of hearing a sound.

<div align="right">(O'Callaghan 2009, p. 609; cf. 2008b, p. 318)</div>

This seems to be the mainstream view, both historically (with support from Aristotle, Berkeley, Strawson, and Warnock) and in the contemporary literature (Nudds 2001, 2010, O'Callaghan 2008b, 2009, Kubovy and Valkenberg 2001, Bregman 1990, Martin 1997).[7] As Mohan Matthen summarizes, "it is the sound of the coach that is directly heard, not the coach itself" (Matthen 2010b, p. 78).

There are very few dissenting accounts. The most famous of these is by none other than Martin Heidegger, who writes in a much-quoted passage:

We never really first perceive a throng of sensations, e.g., tones and noises, in the appearance of things . . . ; rather we hear the storm whistling in the chimney, we hear the three-motored plane, we hear the Mercedes in immediate distinction from the Volkswagen. Much closer to us than all sensations are the things themselves. We hear the door shut in the house and never hear acoustical sensations or even mere sounds. In order to hear a bare sound we have to listen away from things, divert our ear from them, i.e., listen abstractly.

<div align="right">(Heidegger 1977, pp. 151–2)</div>

Although there are some (very few) others who propose, like Heidegger, that we hear objects, not sounds (one such example is Leddington 2013), this view is very often dismissed as being as unusual and eccentric as some of Heidegger's other views.

This debate is deeply intertwined with the one about the ontology of sound (Kulvicki 2008, O'Callaghan 2007, Pasnau 1999, Nudds and O'Callaghan 2009, Casati and Dokic 1994). The big divide in this debate, to simplify things a little, is between those who take sounds to be

---

[7] I do not mean to suggest that these authors present one unified account of auditory individuals. There are various differences between these views in a number of respects—for example, about what they take to be sounds and about whether the source of these sounds is also something that enters the content of our auditory state. I will leave these differences aside.

individuals (Casati and Dokic 1994, O'Callaghan 2007) and those who take them to be qualities of individuals (Kulvicki 2008, Pasnau 1999) (both horns of this dilemma have various versions). If we assume that sounds are individuals—entities, not the qualities of entities—then claiming that the sensory individuals of auditions are sounds would mean that they are not ordinary objects: we attribute properties to sounds perceptually. But if we assume that sounds are qualities of individuals, then sounds could be thought of as exactly those properties that are perceptually attributed to the sensory individuals. So sounds, according to this account, cannot serve as sensory individuals—they are just not the kinds of things that can serve as sensory individuals—they are qualities, not individuals. So those accounts that consider sounds to be the qualities of individuals could be taken to be a version (and in fact, a much more sophisticated version than Heidegger's) of the view that the sensory individuals of audition are objects. Sounds are the properties we perceptually attribute to them.

It is important to note that, strictly speaking, the debate is not about what we hear, but about what we hear "directly" (Matthen 210b, p. 78) or what we hear "in the first instance" (O'Callaghan 2009, p. 609). The proponents of the view that auditory individuals are sounds would allow that we (indirectly, in the second instance) do hear objects. Their proposal is that we only hear objects by virtue of hearing sounds. But this makes the difference between the two alternative views somewhat blurry: it brings in thorny issues about "direct" (as opposed to indirect) perception, for example—a term J. L. Austin famously branded as "a great favourite among philosophers, but actually one of the less conspicuous snakes in the linguistic grass" (Austin 1964, p. 15).

Formulating the question about what we hear in terms of sensory individuals is an attempt at getting rid of the conceptual ambiguity concerning the direct versus indirect objects of perception. Remember, sensory individuals are the entities that our perceptual system attributes properties to. So the general suggestion would be that sensory individuals are what we perceive "directly." There may be some other entities that we perceive "indirectly," but they may not be the same as sensory individuals.[8] As we have seen, some represented properties are attributed

---

[8] Some accounts of auditory individuals may not accept this (maybe Nudds 2001, 2010): they may insist that both what we perceive directly and what we perceive indirectly are (part

non-perceptually, and these may be attributed to an entity that is not the sensory individual. The perceptual (auditory) property-attribution to, say, sounds can of course lead to a non-perceptual property-attribution to objects, but there is only one entity that properties are attributed to perceptually: the sensory individual. If we can decide what the sensory individual of audition is, we can answer the question about what we hear ("directly" and "in the first instance").

3.3.3 Methodological considerations

I will defend a pluralistic view, according to which the sensory individuals of audition are sometimes sounds, sometimes ordinary objects, and sometimes maybe even spatiotemporal regions. And my argument will be provided by the emphasis on pragmatic representations. But first I need to say a word about the methodology used for deciding what the sensory individuals of audition are.

It seems that the most widespread methodology used in this debate is based on intuition and introspection. This is true for both sides of the debate (Heidegger 1977, Leddington 2013, O'Callaghan 2008a, Nudds 2001). As O'Callaghan writes, "auditory experience presents sounds as independent from ordinary material things, in a way that visual and tactual features are not" (O'Callaghan 2008a, p. 804). Matthew Nudds is even more explicit:

[T]he idea that our experience of sounds is of things which are distinct from the world of material objects can seem compelling. All you have to do to confirm it is close your eyes and reflect on the character of your auditory experience.

(Nudds 2001, p. 210)

We have strong reasons to mistrust our intuitions in general (Schwitzgebel 2008, Spener and Bayne 2010, Spener 2011, Swain et al. 2008). Our intuitions depend on random variables like the cleanliness of our surroundings and hands (Schnall et al. 2008, Zhong and Liljenquist 2006), on whether we hold a cup of warm coffee or a glass of cold soft drink (Williams and Bargh 2008), on whether we have just watched an episode of Saturday Night Live or a boring documentary (Valdesolo and

---

of the) sensory individuals. If this is so, then it is even more important to get clear about the nature of the sensory individuals of audition as I defined them: the entity that properties are attributed to auditorily.

DeSteno 2006), on whether we are holding a teddy bear in our hands (Tai et al. 2011), and so on (I will say more about these odd findings in the next chapter). This should make us wary of relying entirely on our intuitions when settling a philosophical debate.

But besides the general unreliability of intuitions, we have more specific reasons to mistrust our intuitions in the case of the sensory individuals of audition. First, perception can be conscious or unconscious, and both conscious and unconscious perceptual processes attribute properties to sensory individuals. When they do so unconsciously, we are simply not in the position to have intuitive or introspective evidence about what the sensory individuals are.

But not everyone uses intuitions or introspection to decide the debate about sensory individuals. Here is Matthen's argument:

> The features or qualities that audition delivers to consciousness are of the following sort: loud, soft, high, low, and so forth. Features of this sort are not attributable to the coach or to its wheels. The squeak of the coach's wheels may be high and the rumble that it makes as it rolls along the road might be low. However, the wheels themselves are not high, and the coach itself is not low.
>
> (Matthen 2010b, pp. 78–9)

Matthen adds immediately after this quote that "This is more than a matter of language" (Matthen 2010b, pp. 78–9), but it is difficult to see what else it is a matter of. Because in our ordinary language linguistic predicates like "being loud" and "being quiet" attach to the subject of "sounds" and not of "ordinary objects," Matthen concludes that our perception attributes the properties that are expressed by these predicates to sounds and not to ordinary objects. But we have no reason to suppose that the way our language describes (mainly unconscious) auditory perception is anywhere close to how auditory perception in fact works.[9]

Instead of relying on evidence from intuitions, introspection, or language, my approach is to use considerations from the connection between perception and action. If perception at least sometimes guides our goal-directed actions, it should carve up the perceived scene in a way

---

[9] There are other arguments that I will not analyze here. A couple of examples: Martin 1997, p. 93, appeals to visual demonstratives; Kubovy and Valkenberg 2001, p. 102, appeal to the figure/ground distinction; and Bregman 1990, p. 2, uses perceptual constancy. O'Callaghan also appeals to some of these considerations (see especially O'Callaghan 2007).

that helps us in performing actions. And this should give us an argument about what the sensory individuals of audition are.

3.3.4 The sensory individuals of pragmatic representations

Pragmatic representations are perceptual states, and some pragmatic representations are states of our auditory perception. In this case, what are the sensory individuals of auditory pragmatic representations? What entities do pragmatic representations attribute properties to?

As the function of pragmatic representations is to facilitate our actions, it seems that they should attribute properties to the parts of the perceived scene in such a way as to facilitate the successful perform-ance of actions with these parts of the perceived scene. But we very rarely perform actions with sounds. Thus, we have reason to suppose that whatever pragmatic representations attribute properties to (directly and in the first instance), they are not sounds. And this casts some doubt on the claim that the sensory individuals of auditory pragmatic repre-sentations are necessarily and exclusively sounds.

What are they then? A tempting answer would be to say that they are ordinary objects: we perform actions with ordinary objects. Thus, pragmatic representations attribute properties to ordinary objects, not to sounds. The proposal then would be that the sensory individuals of auditory pragmatic representations are not sounds: they are objects.

While proceeding this way may seem tempting, we need to be more careful. If pragmatic representations carve up the perceived scene into units that are relevant to one's action, then these units are unlikely to be sounds. Take auditorily guided actions, like jumping away from the quick motorbike approaching from behind. We are not jumping away from the sound of the motorbike, but from the motorbike itself.

Here is another example. Lying in a dark hotel room, I hear a mosquito by my face and I try to kill it without switching on the light. The only way I can attribute action-properties to this mosquito (which I need to do in order to kill it) is by means of auditory perception. In other words, I need to have an auditory pragmatic representation. What does this auditory pragmatic representation attribute action-properties to? Again, a deeply unsatisfying answer would be to say that it attributes properties to the mosquito's sound: I definitely don't want to perform any action directly on the sound the mosquito makes. My auditory

pragmatic representation, it seems, attributes properties to the mosquito itself.

Examples of this kind may establish that the sensory individuals of audition are not (or not exclusively) sounds. But do they establish the positive claim that auditory individuals are ordinary objects? I am not sure. The picture may be more complicated.

We have seen that the trenches in the debate about the sensory individuals of audition are very differently located from the ones in the debate about the sensory individuals of vision. While in the case of vision the debate is about ordinary objects versus spatiotemporal regions, in the case of audition it is about sounds versus ordinary objects. The "spatiotemporal region" view is not even on the radar in the auditory individuals debate. This may be an unfair omission.

Let us go back to the mosquito example. Maybe the auditory individual here is not an ordinary object, that is, the mosquito itself, but the spatiotemporal region it occupies. The suggestion would then be that my auditory pragmatic representation attributes properties to a spatiotemporal region, and it is the attribution of these properties that guides my action of slapping. The action-properties that need to be attributed for the performance of this action could be attributed to a spatiotemporal region as much as they could be attributed to the mosquito. Further, it may seem superfluous to attribute action-properties to the mosquito itself—all that is needed for the performance of the action is the attribution of action-properties to the spatiotemporal region where I'm about to slap.

I do not want to fully endorse this way of analyzing the mosquito example. The main message of this example is that the view according to which sensory individuals are sounds seems mistaken in this case. I want to leave open the question about whether the "ordinary object" view or the "spatiotemporal region" view is better suited to explain this case. But there may be actions where the "spatiotemporal region" view appears more plausible.

Here is an example. For certain hearing tests you are asked to sit in an empty dark room. Your task is to point to the direction where you hear the sounds coming from. Your pointing action in each case is guided by your auditory pragmatic representation. What does this auditory pragmatic representation attribute properties to? Here are the three candidates. First, it seems odd to say that it attributes properties to the sound

you hear, as the pointing action is not performed on the sound—you are not pointing at the sound. Second, it also seems odd to say that it attributes properties to an ordinary object, as you are not aware of any ordinary object around you—remember, you are sitting in a dark empty room. The most natural way of describing what is happening here seems to be that your auditory pragmatic representation attributes properties to a spatiotemporal region, and it is this property attribution to this spatiotemporal region that allows you to perform the action of pointing at this spatiotemporal region. I am not claiming that the "sound" view and the "ordinary object" view could not be tweaked in such a way that it can explain this scenario,[10] but on the face of it this example seems to be a reason to re-evaluate the "spatiotemporal region" view of sensory individuals, at least when it comes to the auditory sense modality.

One may interject that the "spatiotemporal region" view of sensory individuals can be dismissed on independent grounds. As we have seen, there are at least four influential and strong arguments against it: (i) it is too revisionary, (ii) it makes talking about perceptual justification problematic, (iii) it delivers the wrong phenomenology when it comes to perceptually tracking an object through time and space, and (iv) it is in conflict with empirical findings about tracking multiple objects at a given time.

Argument (i) and argument (iii) have the same structure: the "spatiotemporal region" view is in conflict with the way we take ourselves to be perceiving. We take ourselves to be perceiving objects, not places (i), and we take ourselves to be perceiving motion through space, where objects "take their properties with them" (iii). But it is unclear what these considerations have to do with the debate about the nature of sensory individuals, since sensory individuals are not what we take ourselves to be perceiving: they are the entities the properties are perceptually attributed to (consciously or unconsciously). Further, the empirical findings in (iv) concern the visual sense modality only.[11]

---

[10] A good bet for the "sound" view would be the appeal to O'Callaghan's account of what sounds are, where sounds are concrete particulars located around the sound source.
[11] One may take some of the experiments in Bregman's "auditory scene analysis" research program to demonstrate findings similar to (iv) in the auditory sense modality (e.g., Bregman 1990), but the extent of this similarity is not at all clear (see especially Cusack et al. 2004, Wang and Brown 2006).

How about (ii): the argument from perceptual justification? As an analogy, consider the debate about the range of properties that are perceptually attributed. No philosopher claims that the range of properties perceptually attributed is the same as the range of properties attributed per se. I may represent my laptop as having the property of being made in Malaysia in 2010, but this property is very unlikely to be perceptually attributed. Different philosophers draw the line differently: we get a different range of properties that are perceptually attributed, from the minimalistic answer that restricts the range of these properties to shape, size, and color, to a less minimalistic answer that would include sortal properties and maybe even dispositional properties. But the justificatory transition from perception to belief implies that a new range of properties are attributed by a belief that is justified by a perceptual state that does not attribute those properties. Thus, there is a necessary mismatch between the range of properties attributed in perception and the range of properties attributed by belief—no-one seems to worry about the possibility of perceptual justification because of this. And there is no reason to believe that things are different when we consider the debate about the nature of the entities these properties are attributed to. If the mismatch between the perceptually attributed properties and the non-perceptually attributed properties is not worrisome for the prospects of perceptual justification, then the mismatch between the sensory individuals and the particulars our beliefs are about should not be worrisome either. But then the "spatiotemporal region" view may not be completely hopeless as an account of the sensory individuals of audition than it is as an account of the sensory individuals of vision.

It is important not to overestimate the scope of these considerations. I want to refrain from fully endorsing the "spatiotemporal region" view of auditory individuals. But this view seems much more promising in the case of audition than in the case of vision. Hence, it may be a good idea to take the "spatiotemporal region" view of auditory individuals seriously. But the main aim of this section was to argue against the mainstream view that auditory individuals are sounds. Whether we should replace them with ordinary objects or spatiotemporal regions is a further question.

To sum up, I gave an argument that is structurally parallel to the one I outlined in Section 3.2: there we saw that the properties that pragmatic representations attribute should be relevant to one's action. Here the

equivalent claim is that the entities these properties are attributed to should also be relevant to one's action. And the entities that are normally relevant to our actions are not sounds but ordinary objects (or maybe spatiotemporal regions).

This conclusion is not about auditory perception in general. And it is not even about every possible pragmatic representation in the auditory sense modality. My claim was that as we normally perform actions on objects (or spatiotemporal regions), and not on sounds, it seems that pragmatic representations, the function of which is to help us perform actions, attribute action-properties to objects or spatiotemporal regions, not sounds. But it is possible that some people, perhaps professional musicians, do sometimes perform actions with sounds. In these cases we would need to conclude that their pragmatic representations do attribute properties to sounds.

Further, this argument says nothing about auditory perceptual states that are not pragmatic representations, and, presumably, many of our auditory perceptual states are not pragmatic representations. It is possible that some of our non-pragmatic auditory representations attribute properties to sounds—nothing I said here excludes this possibility. But if the argument I present here is correct, then we have strong reason to hold that the sensory individuals of some auditory states are either ordinary objects or spatiotemporal regions. The proponents of the opposite view would need an argument to show that the sensory individuals of some *other* auditory states are sounds.

The conclusion, then, is a pluralistic one: we need to talk about the diversity of auditory objects (this would be a much broader diversity than the one considered in Matthen 2010b). Some auditory individuals are ordinary objects. Some are spatiotemporal regions. And some are sounds.

3.3.5 Sensory individuals in other sense modalities

Can we generalize this argument to other sense modalities? If considerations from pragmatic representations can be used to settle the debate about the sensory individuals of audition, couldn't we use it to settle the debate about the sensory individuals of other sense modalities? I briefly and tentatively consider olfaction.

Bill Lycan gave an argument about olfaction that is in some ways similar to the one I outlined above (Lycan 1996, pp. 147–8, see also Lycan

2000). Lycan is here revising his own earlier theory of olfaction, according to which we smell odors—a claim that would be equivalent to saying that we hear sounds (see also Peacocke 1983, Perkins 1983, Batty 2010, 2011, on this debate). Lycan reconsiders this idea as follows:

> If smells do represent anything, they do after all represent environmental objects of potential adaptive significance. Surely, this is what olfaction is *for*, to signal food, predators, shelter, mates, and other objects of interest ultimately derived from these, and signaling is at least a crude form of representing.
>
> (Lycan 1996, p. 147)

Lycan's final conclusion is that "smells [. . .] represent the environmental entities *by* representing odors" (Lycan 1996, p. 148). But this fails to answer the question about sensory individuals: if smells represent objects by representing smells, then do they attribute properties to objects or to smells? Or to both?

If we can apply the argument I outlined above to olfaction—I'm not fully convinced that we can, partly because of various odd features of this sense modality (see, for example, Porter et al. 2007) and partly because I'm not entirely sure that we need to talk about sensory individuals in the case of olfaction (as Jackson's many properties problem may not even apply)—then the answer we get is that, at least sometimes, olfaction does attribute properties to objects, and not to odors. But, as with the argument about audition, this is not supposed to be a universal claim: there may be instances of olfaction where properties are attributed to odors (for instance, the famous example where one accidentally incinerates a steak, throws out the remains as well as the frying pan, but then finds that the burnt smell still lingers in the kitchen). But the default assumption, as in the case of audition, should be that the sensory individuals of olfaction are ordinary objects (or maybe spatiotemporal regions).

3.4 Pragmatic representations and the dorsal visual subsystem

In following the argument up to this point, many readers have undoubtedly thought about the famous dissociation between the dorsal and ventral visual subsystems. I deliberately refrained from using these labels for describing pragmatic representations. The aim of this section is to clarify the relation between pragmatic representations and dorsal perception.

Humans (and other mammals) have two visual subsystems that use different regions of the central nervous system: the ventral and dorsal streams. To put it simply, the ventral stream is responsible for identification and recognition, whereas the function of the dorsal stream is the visual control of our motor actions. In normal circumstances, these two systems work together, but if one of them is removed or malfunctions, the other can still function relatively well (see Milner and Goodale 1995, Goodale and Milner 2004 for an overview).

If the dorsal stream is malfunctioning, the agent can recognize the objects in front of her, but is incapable of manipulating them or even localizing them in her egocentric space (especially if the perceived object falls outside the agent's fovea). This is called optic ataxia. If the ventral stream is malfunctioning (a condition called visual agnosia), the agent can perform actions with objects in front of her relatively well, but she is incapable of even guessing what these objects are.

The three-dimensional Ebbinghaus illusion I analyzed in the last chapter is normally explained as a nice demonstration of the dissociation between the dorsal and ventral visual subsystems in healthy human adults: the ventral subsystem is fooled by the illusion, but the dorsal is not. The other examples in which optical illusions deceive the eye but not the hand (Ponzo, Müller-Lyer, Kanizsa compression, hollow face, etc.) are analyzed in the same way. Sometimes our ventral visual subsystem attributes a different property to an object from the one the dorsal subsystem does.

In light of these features of the dorsal visual subsystem, a very tempting suggestion would be to say that pragmatic representations are the representations of the dorsal visual subsystem. The dorsal system guides action, as do pragmatic representations. And the dorsal system represents the world in such a way that would help us perform actions, as do pragmatic representations.

The aim of this section is to address the connection between pragmatic representations and the dorsal visual subsystem. My main goal is to carefully detach my claims about pragmatic representations from claims about the dorsal stream. I do not think that pragmatic representations are the representations of the dorsal stream. Before I give four arguments in favor of this claim, I want to point out that the question about the neural implementation of pragmatic representation is a very important one, and I hope that a lot of empirical work will be done on this in the

near future. Whatever the neural implementation of pragmatic representation is, it surely includes the dorsal visual subsystem, but I doubt that it is restricted to the dorsal visual subsystem, for the following four reasons.

First, the anatomical distinction between the dorsal and the ventral visual subsystems is not as neat and clear-cut as was originally thought. It seems that there are interactions between the two streams at various points in perceptual processing (see, for example, Jeannerod 1997, Franz and Gegenfurtner 2008, Franz et al. 2000, Schenk and McIntosh 2010). Further, to make things even more complicated, it has been argued that instead of two visual subsystems we need to talk about three: the ventral, the ventro-dorsal and the dorso-dorsal. To simplify matters considerably, the proposal is that what has been taken to be one single dorsal subsystem should be divided into two: one responsible for manipulating objects (dorso-dorsal), and one responsible for localizing in egocentric space (ventro-dorsal) (Rizzolatti and Matelli 2003). Even more recently, it has been suggested that what was originally taken to be the dorsal stream is in fact the ensemble of three different visual subsystems (Kravitz et al. 2011). To talk about the dorsal stream as an independent chunk of the brain, and to talk about pragmatic representations as the representations of this unified and independent bit of mental processing, would be misleading to say the least.

Second, there is a major debate, both in vision science and in philosophy of cognitive science, about whether dorsal vision is unconscious. The original proposal was that ventral visual processing may be conscious or unconscious, but dorsal processing is always unconscious (see especially Milner and Goodale 1995, Goodale and Milner 2004). But this view has been criticized both on empirical and on conceptual grounds (see, for example, Dehaene et al. 1998, Jeannerod 1997, Jacob and Jeannerod 2003, Clark 2001). This debate does not seem to go away (see Brogaard 2011, forthcoming b, Briscoe 2008, 2009, Milner and Goodale 2008, Jeannerod and Jacob 2005, Jacob 2005, Goodale 2011, Clark 2009, Kravitz et al. 2011). As pragmatic representations can be conscious or unconscious (although they are typically unconscious), if I were to equate pragmatic representations with dorsal perceptual processing I would have to take sides in this grand debate, which I would like to avoid.

The third reason why pragmatic representation should not be identified with the representation of dorsal perception is because of the

multimodality of perception. There is a lot of recent evidence that multimodal perception is the norm and not the exception—our sense modalities interact in a variety of ways (see Spence and Driver 2004, Vroomen et al. 2001, Bertelson and de Gelder 2004 for summaries, and O'Callaghan 2008a, Macpherson 2011 for philosophical overviews). Information in one sense modality can influence the information processing in another sense modality at a very early stage of perceptual processing (even in the primary visual cortex in the case of vision; for example, see Watkins et al. 2006).

A simple example is ventriloquism, which is commonly described as an illusory auditory experience caused by something visible (Bertelson 1999, O'Callaghan 2008b). It is one of the paradigmatic cases of crossmodal illusion: we experience the voices as coming from the dummy, while they in fact come from the ventriloquist. The auditory sense modality identifies the ventriloquist as the source of the voices, while the visual sense modality identifies the dummy. And, as often (although not always—see O'Callaghan 2008b) happens in crossmodal illusions, the visual sense modality wins out: we (auditorily) experience the voices as coming from the dummy. But there are more surprising examples: if you see a flash and you hear two beeps during the flash, you experience it as two flashes (Shams et al. 2000).

Now, as we have seen, pragmatic representations are not necessarily visual: they can occur in any sense modality—in fact, I talked quite a bit in Section 3.3 about auditory pragmatic representations. But the dissociation between the dorsal and ventral subsystem is a distinction in the visual sense modality. Some have suggested a similar dissociation for speech perception (see, for example, Hickock and Poeppel 2007, Kaas and Hackett 1999) and for touch (Reed et al. 2005, Dijkerman and De Haan 2007), but the evidence for dissociations similar to the one in the case of vision is far from clear in audition, olfaction, and the other sense modalities. Tying pragmatic representations to the dorsal visual subsystem would make it difficult to talk about pragmatic representations in sense modalities other than vision.

Further, the literature on the multimodality of perception clearly shows that our perceptual states in one sense modality are influenced by the information we receive in other sense modalities. Since pragmatic representations are perceptual states, we should expect that they are also influenced by information from different sense modalities. And there are

some recent behavioral experiments supporting the multimodality of pragmatic representations (see especially Stein et al. 2004, Gentilucci et al. 1995).

How about the dorsal stream? Although it seems clear that the dorsal stream is also multimodal (see, for example, Battaglia-Mayer and Caminiti 2002), the extent of the crossmodal influences on dorsal processing has been debated (see, for example, Lewis and Van Essen 2000, Rozzi et al. 2008). Again, it seems that the neural correlate of pragmatic representation has a lot to do with the dorsal stream, but the current empirical evidence on multimodal perception does not quite support the claim that it is identical to, or fully exhausted by, the dorsal stream.

Fourth, pragmatic representations are sensitive to various top-down factors: the subject's affective life (Morgado et al. 2011), her language skills (Pulvermuller and Hauk 2005), and her expectations or knowledge. Here is a famous example: two very widely used brands of matches in the UK are "Swan Vestas" and "Scottish Bluebell." The box of Swan Vestas is 25 percent larger than that of Scottish Bluebell. It was tested whether the brand of the matchboxes influences our grip size when grasping them, and it was found that it does (McIntosh and Lashleya 2008). When the subjects were grasping the 1.25-scale replica of the Scottish Bluebell box, their grip size was smaller than it was when grasping the normal Swan Vestas of the same size. And when they were grasping the 0.8-scale replica of the Swan Vestas box, their grip size was larger than it was when grasping the normal Scottish Bluebell box. Hence, the recognition of the brand of the matchboxes influences grip size: it influences pragmatic representations.

But dorsal processing is supposed to be quick and automatic—it is not supposed to be sensitive to top-down influences, especially not factors like the brand of matchbox. Thus, it seems that dorsal vision is not the full story about pragmatic representations.

In short, these findings all point in the direction of a theoretical framework where it is clear that the dorsal stream plays an important role in the implementation of pragmatic representations, but it is unlikely that it plays the only role.

4

Action

All our final decisions are made in a state of mind that is not going to last.

Marcel Proust: *À l'ombre des jeunes filles en fleurs*, 1919.

4.1 Belief–desire psychology and its discontents

Pragmatic representations are part of what constitutes the immediate mental antecedents of action. Further, sometimes they are the only representational states that are needed for the performance of an action. Some other times, in the case of more complex actions, we also have some other representational states, like beliefs and desires. But all actions are triggered (or accompanied) by pragmatic representations.

If this picture is correct, then we have to seriously revise the classic belief–desire model of motivation. As we have seen, according to the classic belief–desire model, it is beliefs and desires that mediate between sensory input and motor output: I see that it is raining outside, I form a belief that it is raining outside, I have a desire not to get wet, I have some further beliefs about the best ways of not getting wet, on the basis of these beliefs and desires I form an intention to take an umbrella, and this leads to the action of taking the umbrella.

I do not deny that in the case of some of our actions we need this mental apparatus to precede our actions—especially in the case of complex actions. But I do want to deny that this picture is the correct one for all of our actions. In fact, in the case of most of our actions (and probably all the actions of non-human animals and small children), there is no need for any beliefs or desires.

This stance is not particularly new. Here is a famous passage by William James that could be taken to make the same point:

Whilst talking I become conscious of a pin on the floor, or of some dust on my sleeve. Without interrupting the conversation I brush away the dust or pick up the pin. I make no express resolve, but the mere perception of the object and the fleeting notion of the act seem of themselves to bring the latter about.

(James 1890, p. 522)

We can flesh out this example to give us the following argument. The actions of picking up the pin or brushing away the dust are brought about without any mediation by beliefs. All that is needed for their performance is "the fleeting notion of the act." And "the fleeting notion of the act" is presumably not a belief. We can and do perform this action without having any beliefs.

The problem with this argument is that while it may or may not establish that *conscious* beliefs are not necessary for action, it remains silent about the possibility that *unconscious* beliefs are necessary for action. The "fleeting notion of the act" could count as a belief, albeit one that is barely conscious or maybe not conscious at all. James's introspective argument (even if we put aside worries about introspective arguments in general, see for example Schwitzgebel 2008) does not work against the view that beliefs (conscious or unconscious) are necessary for action.

Thus, if we want to establish the claim that beliefs are not required for action, and that the only representational state that is required is pragmatic representation, then we need to be able to rule out not only conscious, but also unconscious beliefs.

And here my account of pragmatic representations may help out James's argument. As we have seen in Chapter 2, pragmatic representations are perceptual states: they are not beliefs. So we have found a class of representational states that are necessary for the performance of actions and that are not beliefs (either conscious or unconscious): pragmatic representations.

But we are not yet in the position to conclude that the only representational states that are required for the performance of actions are pragmatic representations. We still need to rule out the following possibility. Let's accept that the representational component of the mental

state that makes actions actions is not a belief: it is some other representation (a pragmatic representation). But it is still a possibility that each time we perform an action we need to have a belief, in addition to our pragmatic representation. The general suggestion then would be this: while pragmatic representation is necessary for action, this does not rule out that beliefs are *also* necessary (Matthen 2005 and forthcoming alludes to this possibility). To rule out this possibility we would need to show that, at least in some cases, the only representational states that are required for performing an action are pragmatic representations.

Let us consider Searle's famous example: while working at my computer, I suddenly, and without planning to do so, jump up from my desk and start pacing up and down (Searle 1983, pp. 84–5). Searle argues that in the case of this action, there is only intention-in-action—neither prior intentions nor beliefs are needed for the performance of this action. While this may sound appealing, should we accept this? Couldn't we resist this claim with the help of the strategy that jeopardized James's argument? In short, couldn't we just say that we do have beliefs and desires when we perform this action, but they are unconscious (maybe because they occur very quickly)?

So the picture would be the following: you are writing at your computer, and very quickly form an unconscious desire to get up and start pacing up and down, as well as an unconscious belief that this desire can be satisfied in such and such a manner. You thereby form an unconscious prior intention, which then leads to an unconscious pragmatic representation that allows you to jump up and start pacing up and down without bumping into the furniture.

The problem with this picture is that it postulates no less than four unconscious mental states, and the only justification for this is to salvage the view according to which beliefs are necessary for action. But, as we have seen, we need to be careful when postulating unconscious representations. We do have very strong reasons for postulating unconscious pragmatic representations, namely that without such representations we would not be able to explain our fine-grained movements (for example, in the three-dimensional Ebbinghaus illusion or in the basketball example). But we have no reason to postulate unconscious beliefs, desires, and prior intention in the Searle example, other than to salvage

the view according to which beliefs are necessary for action. If we don't want to make such an ad-hoc move, then we can conclude that the only representational state in this example that is required for action is the pragmatic representation.

So, to return to the James example, the action of picking up the pin or brushing away the dust could be said to be brought about without any mediation by beliefs or desires. All that is needed for their performance are pragmatic representations. The only properties that need to be attributed for this action to be performed are the action-properties represented by the agent's pragmatic representation. And there are many more examples of such prosaic actions—tying our shoelaces, brushing our teeth, avoiding bumping into people in a crowded underground station, and so on.

The general rhetoric I will follow is to allow that the belief–desire model may be the right way of thinking about complex action, but to also emphasize that most of our actions are very simple and are brought about by pragmatic representations alone, and that even complex actions could not be performed without having pragmatic representations. Even if I do go through all my deliberations about the umbrella, in the end, in order to take the umbrella, I need to represent it as having a certain spatial location that allows me to reach for it, a certain size that allows me to approach it with the right grip size, and as having a certain weight that allows me to exert the right force when lifting it. In short, I need to represent it as having action-properties: I need to have a pragmatic representation. The upshot is that while pragmatic representations are necessary for bringing about an action, beliefs and desires are not.

In Section 4.4, I explore the idea that the belief–desire story is insufficient for complex as well as simple actions.

4.1.1 The cognitive and the conative

We have seen that pragmatic representations are the *representational components* of the immediate mental antecedents of action. They represent all the properties of the object the representation of which is necessary for the performance of the action. But they do not move us to act. Following Myles Brand (1984), I called the component of the immediate mental antecedent of action that moves us to act the

"conative component."[1] And what Brand calls the "cognitive component," I call pragmatic representation.[2]

As the distinction between the cognitive and the conative components of the immediate mental antecedent of action will play an important role throughout the chapter, I need to make some clarifications. First, the claim I make is not that the only mental state that is needed for triggering an action is pragmatic representation. The conative component is also needed—otherwise nothing would move us to act. The claim I make is that the only *representational* state that is needed for bringing about an action is pragmatic representation. We also need the conative component, and just what that may be is a question I would like to avoid here.

We can represent the cup in front of us as having the size-property that allows me to approach it with the right grip size, and having the weight property that allows me to exert the appropriate force when lifting it, but this representation will not in itself bring about the action of lifting the cup. We also need some kind of odd mental state that moves us to act—some kind of trigger: a mental state that may not even have representational content. I said that I will say very little about mental states of this kind, because I do not know how we can talk about the mental state that moves us to act without also solving the problem of free will, and I definitely do not want to do that here.

The distinction between the cognitive and the conative aspects of the immediate mental antecedents of action can also give us a simple way of analyzing those situations where our pragmatic representations do not trigger actions. Here is a nice literary example by Robert Musil:

---

[1] Kent Bach makes a similar (but not identical) distinction between "receptive representation" and "effective representation," which together make up the "executive representation" that is the immediate mental antecedent of action (Bach 1978, see especially p. 366). Some contemporary cognitive scientists also often appeal to a similar duality when discussing motor cognition (see, for example, Haggard 2005a, especially p. 293, and Haggard 2005b).

[2] The James quote I used earlier could be interpreted as being an early version of this distinction. When James says that "I make no express resolve, but the mere perception of the object and the fleeting notion of the act seem of themselves to bring the latter about" (James 1890, p. 522), this sounds like he claims that it takes two factors to bring about an action: "the mere perception of the object" (this would be the cognitive component, the pragmatic representation) and "the fleeting notion of the act" (which, at least on some interpretations, would play the role of the conative component).

I have never caught myself in the act of willing. It was always the case that I saw only the thought—for example when I'm lying on one side in bed: now you ought to turn yourself over. This thought goes marching on in a state of complete equality with a whole set of other ones: for example, your foot is starting to feel stiff, the pillow is getting hot, etc. It is still a proper act of reflection; but it is still far from breaking out into a deed. On the contrary, I confirm with a certain consternation that, despite these thoughts, I still haven't turned over. As I admonish myself that I ought to do so and see that this does not happen, something akin to depression takes possession of me, albeit a depression that is at once scornful and resigned. And then, all of a sudden, and always in an unguarded moment, I turn over. As I do so, the first thing that I am conscious of is the movement as it is actually being performed, and frequently a memory that this started out from some part of the body or other, from the feet, for example, that moved a little, or were unconsciously shifted, from where they had been lying, and that they then drew all the rest after them.[3]

This happens to us all the time. The pragmatic representation that allows us to perform the action is there all along, but the action is just not coming. This phenomenon is very easy to explain if we make a distinction between the cognitive and the conative components of the mental antecedents of our actions: we only move when the conative aspect is also there. The cognitive aspect—the pragmatic representation—does not, in itself, lead to any action.

Another important consequence of the cognitive/conative distinction concerns the psychological disorder called utilization behavior. Utilization behavior (Lhermitte 1983, Shallice et al. 1989) is caused by lesions in the frontal lobe. Patients with utilization behavior tend to perform actions with the objects they see, regardless of whether they need to perform this action or whether it is socially acceptable to do so. For example, they climb into a bed if they see one, even if it is the middle of the day; if they notice a pair of glasses in front of them, they put them on, even if they do not normally wear glasses and even if they are already wearing a pair; they open an umbrella even if it is not raining or if they are indoors, etc.

One way of analyzing utilization behavior is to posit a mental state in healthy humans, the function of which is to suppress the automatic performance of actions (see, for example, Pacherie 2000, 2007, Frith

---

[3] Robert Musil: *Diaries*. New York: Basic Books, 1999, p. 101. See also Goldie 2004, pp. 97–8.

et al. 2000). The neural correlates of this mental state are taken to be located in the frontal lobe, and that is why, in patients with utilization behavior whose frontal lobe is damaged, this mental state is missing. And that is why they have trouble suppressing their actions. According to this picture of utilization behavior then, in healthy humans the immediate mental antecedent of action is intrinsically motivating. What can stop it from being intrinsically motivating is an extra, suppressing, mental state. The presence of this extra mental state is phylogenetically relatively new, and it is exactly this that is missing in patients with utilization behavior.

But the distinction between the cognitive and the conative components of the immediate mental antecedents of action can help us explain utilization behavior in a much simpler manner—without positing, in a more or less ad-hoc manner, an extra mental state. According to this alternative picture, utilization behavior results from the malfunctioning of the conative component of the immediate mental antecedent of action. The cognitive component—the pragmatic representation—is unimpaired: the patients perform these actions with impressive success, the right grip size, etc. But whenever they form a pragmatic representation, this automatically triggers the action.

Thus, the difference between healthy humans and patients with utilization behavior is not a difference external to the immediate mental antecedent of action. It is a difference in one of the two components of the immediate mental antecedents of action: the conative one. And this difference is not a matter of the impairment of the mechanism that would suppress actions, but a matter of the oversensitivity of the mechanism that would move the agent to act.[4]

An obvious advantage of this alternative picture is that it does not force us to posit an additional, suppressing, mental state. Another

---

[4] Note that the proponents of explaining utilization behavior in terms of the lack of a suppressing mechanism could still endorse the cognitive/conative distinction. They claim that the immediate mental antecedents of action are intrinsically motivating. I agree. I further argue that these intrinsically motivating immediate mental antecedents of action consist of two separate components, the cognitive and the conative ones. And I see no reason why the proponents of explaining utilization behavior in terms of the lack of a suppressing mechanism could not make the same distinction. But if they do so, then they would have all the means to explain what works differently in the case of utilization behavior patients—namely, the conative component. There is no need to postulate a further, external, suppressing mental state.

advantage is that it seems to fit the empirical data about utilization behavior better. Those regions of the temporal lobe the lesions of which cause utilization behavior are the mesial frontal lobe and the frontostriatal pathways (see Archibald et al. 2001). While these regions were traditionally associated with action suppression (e.g., Hess et al. 1946), it is now agreed that they are associated with action initiation and organization, not primarily inhibition (see, for example, Jeannerod 1997, pp. 151–2, Shallice and Burgess 1991).

A third reason why we may want to make a distinction between the cognitive and the conative components of the immediate mental antecedents of action is that if we don't, then it is difficult to see how we can avoid the conclusion that pragmatic representations are intrinsically motivating. According to my account, the immediate mental antecedents of action are intrinsically motivating, but only because they include the conative component. Pragmatic representations are not intrinsically motivating: we can represent objects as having action-properties without being moved to act.

But if we don't allow for the cognitive/conative distinction, then it would seem to follow that pragmatic representations themselves are intrinsically motivating—as there is no separate (and optional) conative component that would accompany them. This would lead to one way of cashing out the general suggestion I mentioned in Chapter 2 about the "double direction of fit" of pragmatic representations. According to this way of thinking about pragmatic representations, they would have both a descriptive and a prescriptive direction of fit: by representing the properties of the objects in front of us, they also move us to act—they are intrinsically motivating.

I argued against the general suggestion of the "double direction of fit" in Chapter 3. According to my account, pragmatic representations do not have a double (both descriptive and prescriptive) direction of fit. They only "describe"—represent objects as having action-properties. Pragmatic representations allow us to perform actions, but they do not automatically trigger actions.

To sum up, we have good reason to make a distinction between the representational (cognitive) and conative components of the immediate mental antecedents of action. Pragmatic representations are the former. This distinction will play an important role later in this chapter.

4.2 Naturalizing action theory

About 30 years ago, a number of philosophers of action were urging a naturalist turn in action theory. This turn did not happen. My aim is to argue that if we accept the argument about the centrality of pragmatic representations in bringing about actions, we have strong reasons to naturalize action theory.

The most important proponent of the naturalization of action theory was Myles Brand. Brand (1984) argued that philosophy of action should enter its "third stage" (the first one was in the 1950s and 1960s, the second in the 1970s), the main mark of which would be its continuity with the empirical sciences. Brand's methodology for philosophy of action is a package deal. He endorses the following three guidelines for the methodology that action theorists should follow:

(a) Philosophy of action should be continuous with the empirical sciences
(b) Philosophy of action should not privilege intentional actions
(c) Philosophy of action should be independent from ethics/moral philosophy

The last 30 years of philosophy of action could be described as doing the exact opposite of what Brand suggested. Contemporary philosophy of action is almost entirely about intentional actions (and not actions in general), and it is far from being independent from ethics/moral philosophy: in fact, it has (with some rare exceptions) virtually become part of ethics/moral philosophy. Most importantly, contemporary philosophy of action is not, generally speaking, a naturalist enterprise: it consistently ignores empirical findings about actions and its mental antecedents, and has no patience for the cognitive neuroscience of action, for example.[5]

Interestingly, however, a similar naturalist turn (or at least half-turn) did occur in contemporary philosophy of perception. More and more contemporary philosophers of perception seem to have very similar

---

[5] A notable exception is the recent philosophical literature on the "illusion of free will": the sense of agency and conscious will (see, for example, Libet 1985, Wegner 2002, Haggard and Clark 2003, Pacherie 2007). It is important to acknowledge that experimental philosophers do use empirical data in our intuitions about actions and our way of talking about them. But even experimental philosophers of action tend to ignore empirical findings about action itself (as opposed to our intuitions about it).

methodological commitments as the ones enumerated above (see also Nanay 2010d):

(a) Contemporary philosophy of perception takes empirical vision science very seriously
(b) Contemporary philosophy of perception tends not to privilege conscious perception
(c) Contemporary philosophy of perception tends to be independent from epistemology

In recent years, paying close attention to empirical findings about perception seems to be the norm, rather than the exception. What this means is not that philosophy of perception has become theoretical vision science. Rather, philosophical arguments about perception are constrained by, and sometimes supported by, empirical evidence. Even in the case of some of the most genuinely philosophical debates, such as the representationalism versus relationalism debate, many of the arguments use empirical findings as premises (see, for example, Pautz 2010, Nanay forthcoming c). And the fact that many of these empirical findings are about non-conscious perceptual processes shifts the emphasis away from conscious perceptual experience.

Epistemology has always had special ties to philosophy of perception, traditionally because of the role perception is supposed to play in justification. But in contemporary philosophy of perception, perception is no longer only interesting inasmuch as it can tell us something about knowledge. Quite the contrary: epistemological considerations are often used to answer intrinsically interesting questions about perception.[6]

The general picture that these methodological commitments outline is one where philosophy of perception is an autonomous field of philosophy that has important ties to other fields but does not depend on them, and that is sensitive to the empirical findings of vision science. This is

---

[6] One important example comes from Fred Dretske's work. The original link between perception and knowledge is at least partly due to the works of Fred Dretske over the decades (starting with Dretske 1969). Dretske's recent writings, however, turn the established connection between perception and knowledge on its head. He is interested in what we perceive, and some of the considerations he uses in order to answer this question are about what we know (see Dretske 2007, 2010). Dretske's work exemplifies a more general shift of emphasis in contemporary philosophy of perception.

very similar to the picture that Brand envisaged for philosophy of action, but that never in fact materialized.

My aim is to argue that since pragmatic representations are not normally accessible to introspection, naturalized action theory is the only plausible option. Philosophy of action should turn towards philosophy of perception for some methodological support (see also Nanay forthcoming d). As pragmatic representations are both perceptual states and the representational components of the immediate mental antecedents of action, it is the joint job of philosophy of action and philosophy of perception to characterize them. I will argue that this can only be done by relying on the empirical sciences.

4.2.1 Naturalism about action theory

I need to be explicit about what I take to be naturalism about action theory. I have been talking about sensitivity to empirical results, but this is only part of what naturalism entails. The most important naturalist slogan since Quine has been the continuity between science and philosophy. As Quine says,

I admit to naturalism and even glory in it. This means banishing the dream of a first philosophy and pursuing philosophy rather as a part of one's system of the world, continuous with the rest of science.

(Quine 1984, pp. 430–1)

Naturalism in the context of philosophy of action can be, and has been, formulated in a similar manner. Brand, for example, talks about "the integration of the philosophical with the scientific" (Brand 1984, p. x).

Just what this "continuity" or "integration" is supposed to mean, however, remains unclear. More specifically, what happens if what science tells us is in conflict with what folk psychology tells us? Brand clearly hands the decisive vote to folk psychology. As he says, "Scientific psychology is not free to develop any arbitrary conceptual scheme; it is constrained by the conceptual base of folk psychology" (Brand 1984, p. 239). But that has little to do with naturalism, as Slezak (1987, 1989) points out (see especially the detailed point-by-point analysis of how Brand's theory fails on its own terms in Slezak 1989, pp. 140–1, 161–3). If the only role science is supposed to play in action theory is to fill in the details of the pre-existent, unchangeable conceptual framework of folk psychology, then science is not playing a very interesting role at all—the conceptual framework of action theory would still be provided by folk psychology.

Brand's theory, in spite of its false advertisement, is not naturalistic in any sense of the term that would do justice to the Quinean slogan.

What would then constitute a naturalized action theory? We can use Brand's original formulation as a starting point: naturalized action theory urges the integration of the philosophical with the scientific, but a very specific kind of integration: one where the philosophical does not automatically trump the scientific. If it turns out that some of our key folk psychological concepts in philosophy of action (like those of "action" or "intention") fail to pick out any natural kinds, we have to replace them with concepts that do pick out natural kinds.[7] And science can tell us what this new concept should be.

I talked about the importance of empirical findings in naturalized action theory: empirical findings constrain the philosophical theories of action we can plausibly hold. But the interaction between philosophy and the empirical sciences is bidirectional. The philosophical hypotheses and theories, as a result of being empirically informed, should be specific enough to be falsified or verified by further empirical studies. Psychologists and neuroscientists often accuse philosophers in general, and philosophers of mind in particular, of providing theories that are too general and abstract: that are of no use for the empirical sciences.

Philosophers of a non-naturalistic creed are of course free to do so, but if we want to preserve the naturalistic insight that philosophy should be continuous with the empirical sciences, such a disconnect would not be permissible. Thus, naturalistic philosophy needs to give exact, testable hypotheses that psychologists as well as cognitive neuroscientists of action can engage with. Naturalized action theory, besides using empirical studies, could also be used for future empirical research. This is the only sense in which the "integration of the philosophical with the scientific" that Brand talked about does not become a mere slogan. And this is the methodology that has been used by more and more philosophers of perception (I won't pretend that it has been used by all of them), and, given the extremely rich body of empirical research (especially in

---

[7] I am using here the widely accepted way of referring to natural kinds as the real joints of nature because it is a convenient rhetorical device, but I have my reservations about the very concept, for a variety of reasons (see Nanay 2009b, 2010b, 2010h, 2010i, 2011f, 2011g, 2012c, 2012f).

the cognitive neuroscience of action),[8] more and more philosophers of action should use the same methodology.

This may sound like a manifesto about how nice naturalized action theory would be, but the aim of this section is to argue that it is difficult to see how naturalized action theory can be avoided.

4.2.2 Pragmatic representations are not normally accessible to introspection

Remember the argument about perceptual learning, where our pragmatic representation changes, but our perceptual experience does not. Or remember the three-dimensional Ebbinghaus illusion. We have two different mental states in these scenarios: a conscious, incorrect one, and a pragmatic representation, which is (more or less) correct. They are both representations: they both attribute properties to the same object. But they attribute different properties. The conscious experience attributes the size-property we experience the chip as having, and the pragmatic representation attributes the size-property that guides our (successful) action.

Importantly—given that we have a conscious and incorrect representation at the same time as we have a (more or less) correct pragmatic representation of the same properties of the same object— this pragmatic representation must be unconscious. Our conscious perceptual experience attributes a certain size-property to the chip, but our pragmatic representation attributes another size-property—it can only do so unconsciously. Hence, pragmatic representations are (normally) unconscious.

We need to be careful about what is meant by unconscious here. Do these states lack phenomenal consciousness or access-consciousness (Block 1995)? Is it visual awareness or visual attention that is missing (Lamme 2003)? Luckily, we do not have to engage with the Byzantine details of these distinctions (but see Prinz 2010). What matters for the purposes of my argument is that pragmatic representations are not accessible to introspection. When we are grasping the chips in the three-dimensional Ebbinghaus scenario, we have no introspective access to the representation that guides our action and that represents the size of

---

[8] The literature is too large to survey, but an important and philosophically sensitive example is Jeannerod 1997.

the chip (more or less) correctly. We have only introspective access to the conscious perceptual experience that represents the size of the chip incorrectly. Pragmatic representations are not normally accessible to introspection.

A final objection. I said that pragmatic representations are not normally accessible to introspection. But am I justified to use the word "normally" here? Couldn't one argue that the scenarios I analyzed are the "abnormal" ones? I don't think so. Here is a so far unmentioned body of empirical evidence that demonstrates this. If the location (or some other relevant property) of the target of our reaching or grasping actions suddenly changes, the trajectory and/or velocity of our movement changes very quickly (in less than 100 milliseconds) afterwards. The change in our movement is unconscious: subjects do not notice this change, and as it occurs within 100 milliseconds of the change in the target's location, this time is not enough for the information to reach consciousness (Paulignan et al. 1991, Pelisson et al. 1986, Goodale et al. 1986, see also Brogaard 2011). In short, the subject's pragmatic representation changes as the target's location changes, but this change is not available to introspection. And this is true of all actions that require microadjustments to our ongoing action, which means it is true of most of our perceptually guided actions (see also Schnall et al. 2010 for some further structurally similar cases).

If the argument I present here is correct, then pragmatic representations are not normally accessible to introspection. Now we can use this argument to conclude the necessity of naturalizing action theory.

If we accept that pragmatic representations are not normally accessible to introspection, then we have a straightforward argument for the need to naturalize action theory. If the representational component of the immediate mental antecedent of action is not normally available to introspection, then introspection obviously cannot deliver any reliable evidence about it.[9]

Introspection, of course, may not be the only alternative to scientific evidence. There may be other genuinely philosophical ways in which we

[9] It could be argued that the other, conative, "moving to act" component of the immediate mental antecedent of action is also normally inaccessible to introspection, which would further strengthen the case for a naturalized action theory. See Nanay forthcoming d.

can acquire information about a mental state: folk psychology, ordinary language analysis, conceptual analysis, etc. But note that none of these philosophical methods are in a position to say much about pragmatic representations. Pragmatic representations are not part of our folk psychology—as we have seen. When we think about other people's mental states, we think about their beliefs, desires, and wishes, and not so much about the ways in which their perceptual system represents the shape properties of the objects in front of them. Similarly, talk about pragmatic representations is not part of our ordinary language—ordinary language analysis will not get us far. How about conceptual analysis? Arguably, the generation of action theorists that gave us the distinction between the cognitive and conative components of the imme-diate mental antecedents of action (Brand 1984, Bach 1978) did use conceptual analysis, or, more precisely, some version of a transcendental argument: we need to postulate this distinction in order to explain a number of odd features of our behavior.

I see nothing wrong with this approach, but it has its limits. We can, and should, postulate certain mental states—more specifically, pragmatic representations—in order to be able to explain some features of our goal-directed actions, but postulating is only the first step. The real work is in figuring out what these representations are, what properties they represent objects as having, how they interact or fail to interact with the rest of our mind, and so on. And this is something that conceptual analysis is unlikely to be able to do.

Hence, it seems that the only way to find out more about pragmatic representations is by means of empirical research. We have no other option but to turn to the empirical sciences if we want to characterize and analyze them. And as pragmatic representations are the representational compon-ents of what makes actions actions, this means that we have no other option but to turn to the empirical sciences if we want to understand what actions are. Relying on empirical evidence is not a nice, optional feature of action theory: it is the only way action theory can proceed.

4.3 Semi-actions

Some bodily movements are actions. If I decide to eat some yoghurt and get up from my computer to do so, I perform an action. If my leg moves

because the doctor is examining my knee-jerk reflex with her hammer, I do not perform an action—it is a mere bodily movement. As we have seen, one of the most important tasks of action theory is to tell us the difference between actions and mere bodily movements: to tell us what makes actions actions.

In the case of these two examples, it is clear enough whether we should count them as actions or mere bodily movements. If an account of what makes actions actions classifies knee-jerk reflexes as genuine actions, we have reason to be suspicious. But, there are potential examples that do not fall so clearly into one of these two categories: these examples that are neither full-fledged actions nor mere bodily movements I call "semi-actions."

Why are semi-actions important? Philosophers of action assume that there is a mental state type, call it M, that makes actions actions—that constitutes the difference between a mere bodily movement and an action. If our bodily movement is triggered (or maybe accompanied) by M, it is an action. If it is not, it is not an action. There doesn't seem to be any logical space for a middle way in between. It is difficult to see how there can be bodily movements that we characterize as somewhere between mere bodily movements and full-fledged actions: as "semi-actions." Hence, if there are semi-actions, we may need to revise this picture of what makes actions actions.

This is also a potential problem for my account, as it is for any account of what makes actions actions. But the reason why I discuss the possibility of semi-actions at length is not merely as a potential objection that my account faces. I argued for a naturalized philosophy of action in Section 4.2. A key feature of naturalized philosophy of action, as we have seen, is that considerations from everyday folk psychology or ordinary language terms do not automatically trump other considerations. Consider the possibility of semi-actions—a phenomenon that is not part of our everyday folk psychology or our ordinary language— which provide a good test case for just how naturalistic my approach is.

Consider the following six examples:

(a) I decide to eat some yoghurt and get up from my computer to do so.
(b) While typing at my computer, I suddenly and without planning to do so jump up and start pacing around the room (Searle 1983).

(c) I stand behind a strong piece of plexiglass knowing that there is plexiglass in front of me, and when someone on the other side of the glass throws a beach ball at me I reach out in an attempt to catch the ball (the example is from Nanay 2012a).

(d) "I put my face close to the thick glass-plate in front of a puff-adder in the Zoological Gardens, with the firm determination of not starting back if the snake struck at me; but, as soon as the blow was struck, my resolution went for nothing, and I jumped a yard or two backwards with astonishing rapidity" (Darwin 1899, p. 18).

(e) Anarchic hand syndrome patients exhibit clearly goal-directed, well-executed, and complex movements of an upper limb that are nonetheless unintended and sometimes conflict with the agent's desires or even the action they intentionally perform with their other hand (Della Sala et al. 1991, Giovannetti et al. 2005, see also Pacherie 2007 for a philosophical summary).

(f) The doctor is examining my knee-jerk reflex with a hammer and my leg moves as a result.

We have already seen that (a) and (b) are clear examples of action. In fact, case (a) is an *intentional* action (thankfully, in this context I can ignore what that means). Case (b) is also an action, although opinions differ about whether it is an intentional action. Case (f) is a mere bodily movement—no action is performed.

The problem is with the remaining three examples that seem to be somewhere between action and mere bodily movements. And, as we have seen, if they really are semi-actions then the standard account of what makes actions actions needs to be revised. One could, of course, object that these cases are not semi-actions at all: intuitions say that these really are genuine actions (or that they are mere bodily movements). Hence, they pose no challenge to the standard picture of what makes actions actions. I find this way of responding to the challenge too cheap. The threat semi-actions pose to any account of what makes actions actions in general, and my own account in particular, needs to be taken seriously and should not be dismissed on the basis of (unreliable) intuitions.

It can be demonstrated that while people have no problem categorizing cases like (a) and (b) as actions, and (f) as non-action, they are torn when they are asked about cases like (c), (d), and (e) (Nanay forthcoming g).

While some people may have the intuition that (c), (d), and (e) are genuine actions (or that they are mere bodily movements), these people are the minority: most people characterize these cases as "somewhere in between": as semi-actions. Even if your intuition tells you that there are no semi-actions, nothing follows from this: the intuitions of most people tell them that there are. Thus, we need to take the idea of semi-action seriously.

4.3.1 Explaining semi-actions

We have a puzzle then. According to the standard picture, if the bodily movement is triggered (or maybe accompanied) by a specific kind of mental state, it is an action. If it is not, it is a mere bodily movement. As we have seen, this way of formulating the question makes it difficult to explain the intermediary cases like (c), (d), and (e). The bodily movement is either triggered (or maybe accompanied) by this kind of mental state or it is not. It is difficult to see what would constitute half-triggering a bodily movement.

But we can solve this puzzle. In order to do so, we need to go back to the distinction between the cognitive and the conative components of the immediate mental antecedents of action. As we have seen, the general insight is that the immediate mental antecedent of action has two distinct components: one that represents the world in a certain way, and the other that moves us to act.

According to Myles Brand, who introduced the distinction (as well as for others who made similar distinctions, like Bach 1978), both components are needed for a bodily movement to count as an action (see especially Brand 1984, p. 45). This would give a dogmatic answer to our question about the problem cases of action attribution. For Brand, there is no logical space between actions and mere bodily movements: there is no logical space for semi-actions. But this distinction between the two distinct components of the immediate mental antecedents of action could help us to give the outlines of a non-dogmatic answer.

As long as we acknowledge that the immediate mental antecedents of action consist of two distinct components, we can give a coherent account of the intermediary cases of action attribution. The upshot is simple: typically, in the case of performing actions both components are present, and in the case of mere bodily movements neither of them are.

In the case of semi-actions, only one of the two components is present. The component that represents the world in a certain way is present, while the component that is supposed to move us to act is missing. But let us proceed more slowly.

In (a) and (b), the immediate mental antecedent of action has both a cognitive and a conative component. The conative component in (b) is unplanned but nonetheless it is still there: what moves me to act is not something external, it comes from within (although it may not be based on a previously established plan or "prior intention"). And the cognitive component is also there: in (a) I need to represent the contents of the fridge (as well as the whereabouts thereof), and in (b) I need to have some kind of representation of where in my office one can pace up and down and when to turn around in order not to bump into the bookshelf. In the case of (f), in contrast, we have neither of these two components of the immediate mental antecedent of action: neither the cognitive nor the conative one—we don't need to represent the world in any way in order for our knee-jerk reflex to work, and what moves us to act is not the conative component of the immediate mental antecedent of action but the doctor's hammer.

In (c), (d), and (e), the conative component is missing, while the cognitive component is present. We are "moved to act" by some external stimulus. This is what makes (c), (d), and (e) somewhat similar to reflex movements, where there is no conative component either. But, importantly, the cognitive component is there: the cognitive component of our immediate intention is guiding our ongoing activity. That is why we reach out for the ball in (c) in a certain way that reflects the size of the ball and the direction it is coming from. And that is why Darwin jumps in a certain direction (and not others) in (d). Finally, this is the reason why anarchic hand patients can perform extremely complex goal-directed movements with their anarchic hand, and do so successfully, in (e).

This is an important contrast between what I call "semi-actions" and what Brian O'Shaugnessy called "subintentional acts" (O'Shaugnessy 1980, pp. 58–73). Tapping one's foot when listening to music is a subintentional act, but it is not a semi-action—you do not need to represent the world in any way in order to be able to tap your foot, whereas, as we have just seen, you do need to represent the world in a certain way in order to perform semi-actions.

One way of understanding the difference between what I call semi-actions and what O'Shaugnessy calls subintentional acts is as follows. In the case of semi-actions, the cognitive component of the immediate mental antecedent of action is present, but the conative is missing. On the other hand, subintentional acts could be interpreted as acts where the cognitive component is missing, while the conative is present: what moves me to tap my foot when listening to music is not something external—it comes from within.

Some other potential examples (that O'Shaugnessy may or may not consider to be bona fide subintentional acts) include: blinking (without planning to do so, and without any external influence that would trigger the blinking reflex) and swallowing (without planning to do so, and without any external influence that would trigger the swallowing reflex). In these cases, we do not need to represent anything in order for an action to be performed. So the cognitive component is missing: there is no pragmatic representation here. But there is a conative component: whatever moves me to act is not something external (as in the blinking/ swallowing reflex case), but something that comes from within. Do these actions count as bona fide actions? Do they count as subintentional acts? I am not sure. What is more important than the labels is that the distinction between the cognitive and the conative components of the mental antecedents of action allows us to differentiate these importantly and interestingly different cases.

Note that the argument I gave for accommodating the possibility of semi-actions would work for any account of what makes actions actions that allows for the distinction between the cognitive and the conative components of the immediate mental antecedents of action. As my account also allows for this distinction, my account can also accommodate the possibility of semi-actions.

4.4 The last refuge of the belief–desire model

The general rhetoric of this chapter so far has been the following: most of our actions are similar to animal actions. As often the only representational state needed for the performance of these actions are pragmatic representations, it is unlikely that the belief–desire model is the best way of describing them. Therefore, the belief–desire model needs to be

discarded as a general model for what mediates between sensory input and motor output. But the belief–desire model is very useful when we are trying to explain some complex, uniquely human, deliberative actions. In other words, rather than discarding the belief–desire model altogether, we should limit its scope to highly deliberative human actions—for simplicity, I call these highly deliberative human actions "decision-making" in what follows. The aim of this last section is to attack this last refuge of the belief–desire model.

Before I begin, it is very important to clarify what I mean by the belief–desire model in this more specific context of explaining decision-making or highly deliberative human actions. The way I will use this concept is the way it is used in philosophy of mind—not a specific psychological model, but rather a "scientific research program" that aims to explain our decision-making in terms of beliefs and desires. Imre Lakatos introduced the concept of "scientific research program," by which he meant not a specific theory, but a temporal sequence of a set of theories that have a common conceptual or theoretical core (Lakatos 1970, 1974). When I talk about the belief–desire model of decision-making, what I mean is a scientific research program, rather than a specific psychological model that would be one of many in this scientific research program.

There are, of course, many specific psychological models that belong to the scientific research program that I call the "belief–desire model." Some of these are better off than others—I analyze some of these differences below. But my aim is not to criticize some specific psychological models of decision-making—some specific theories belonging to this scientific research program. The aim is to find an alternative to the general scientific research program that aims to explain decision-making in terms of beliefs and desires. And this alternative (again, not an alternative specific psychological model, but an alternative scientific research program) aims to explain decision-making in terms of imagination.

4.4.1 Odd facts about decision-making

Decision-making is not a monolithic category: we can call any action "decision-making" that involves choosing between two potential actions and that involves some deliberation. These two potential actions can be checking two different boxes in a questionnaire, or pushing one button

rather than another. They can also be decisions to marry one person rather than another, or taking up one job offer rather than another.

How does the belief–desire model describe the decision-making process? Here is a somewhat simplified account: the agent has some desires (or other pro-attitudes, such as preferences) and some background beliefs, such that deciding between two possible actions is a matter of comparing the satisfaction of these desires given the background beliefs in the case of the performance of each action. Most versions of the belief–desire model allow for this comparison (or for the beliefs/desires involved in it) to be non-conscious and non-explicit. But they are all committed to the claim that this comparison is about beliefs and desires.

Both the classic rational choice theory literature and its main alternative, the prospect theory, use this general framework, despite the various differences between them. In both cases, the decision is the outcome of a mental process involving beliefs and desires (this is an explicit assumption of rational choice theory, but it is also an implicit—and sometimes explicit—premise in prospect theory, see especially Fox and Tversky 1998, Heath and Tversky 1991, Fox and See 2003, Wakker 2004).

The problem with this way of describing decision-making is that there is a wealth of recent empirical findings about how we actually make decisions that seem to flatly contradict this picture: our actual decision-making is sensitive to order effects, to framing effects, and even to such banal environmental factors as the dirtiness of one's hands. Maybe the way we *should* make decisions is by comparing the satisfaction of our desires, given our background beliefs in the case of the performance of each action, but it is unlikely that this is what in fact happens when we make decisions. In fact, as a consequence of these results, the psychological and decision science literature moved from questions about rational decision-making (what makes decision-making rational?) to questions about our *actual* decision-making (the *locus classicus* is Kahneman and Tversky 1979, see also Bell et al. 1988 and Yaari 1987 on the differences between these two projects). I am engaging with this latter project: the explanation of how we make decisions. Nothing I say here counts against the view that *rational* decision-making is a property described by the belief–desire model (I will come back to this distinction below).

A couple of famous examples. It has been shown that the wording of the task influences decision-making: depending on whether the same

scenario is formulated in terms of the word "kill" or the word "save," subjects make different decisions (Petrinovich and O'Neill 1996). The order in which the questions are raised also influences one's decision-making: the subject's decision will be different depending on whether task A comes before task B, or vice versa (Swain et al. 2008, Petrinovich and O'Neill 1996).

Further, the way the question or problem is framed also influences our decision-making: if the subject has to decide between a certain amount of money or a 50 percent chance of twice that amount of money, the decision depends on whether the subject is about to lose or gain this sum (Tversky and Kahneman 1981). It is the same decision problem, but depending on whether it is framed as gain or loss, the outcome of the decision-making is significantly different.

But decision-making is also influenced by what—completely unrelated—experiences the subject had right before the decision-making process. As Valdesolo and DeSteno (2006) demonstrated, watching an episode of Saturday Night Live or a (boring) documentary on a small Spanish village has a significant impact on the decision made afterwards.

Perhaps most surprisingly, such prosaic factors as the cleanliness of one's environment, and even of one's own hands, also influences one's decision-making. Subjects make significantly different decisions (for example, moral judgments) depending on whether they are seated in a dirty cubicle with leftover pizza boxes and greasy stains on their desk, or in a clean environment (Schnall et al. 2008). Hand washing or cleansing also influences our decisions about what we would be willing to do in specific situations (Zhong and Liljenquist 2006).

It has also been pointed out that whether we are holding a cup of warm beverage or a glass of cold drink also influences our decisions about other people (Williams and Bargh 2008). I need to emphasize that these are not isolated examples for external effects influencing our decision-making— here is one additional, quite evocative, example: there is evidence that holding a teddy bear (as opposed to merely seeing one) influences one's decisions about social behavior (Tai et al. 2011).[10]

---

[10] This is by no means a complete list of the empirical findings that are difficult to explain within the framework of the belief–desire model. Here are four further sets of findings: (a) *Decision-making under risk in non-human animals*: there has been a lot of

Is it possible to accommodate these findings within the framework of the belief–desire model? Remember, according to the belief–desire model, our decision is based on comparing the satisfaction of our desires given our background beliefs in the case of the performance of each action. If this were the case, it is difficult to see how the findings enumerated above would be possible: our background beliefs and desires are not altered by the mess on the table in front of us, or by the film sketch we have just watched.

One may worry that each of these empirical findings only demonstrates that some external factors influence some specific kinds of decision-making: holding a cup of warm coffee influences our decision-making about the personality-traits of others, but not any other kinds of decision-making processes. Similarly, whether the desk in front of us has pizza stains on it influences our decisions about moral evaluations, but not other kinds of decisions. How much can we generalize from these findings then?

The answer is that these findings, put together, demonstrate that decision-making is subject to a wide variety of influences that would not be there if it were a matter of calculating beliefs and desires. Whether there are sub-categories of decision-making that are immune to these influences is an open question—there may be some: for example, decisions concerning mathematical or logical tasks (but see

research on decision-making under risk in non-human animals (Caraco 1981, Bateson 2002, Chen et al. 2006, Hurly and Ossen 1999, Marsh and Kacelnik 2002, McCoy and Platt 2005, Lakshminarayanan et al. 2011). If one tried to explain these risk-biases in terms of the belief–desire model of decision-making (i.e., in terms of comparisons between the satisfaction of desires given background beliefs), this would lead to a serious over-intellectualizing of the minds of very simple animals. (b) *Cross-cultural perspectives on decision-making*: there is a growing body of literature on cross-cultural differences in decision-making (Hsee and Weber 1999, Lau and Ranyard 2005). The majority of these differences are difficult to explain merely in terms of beliefs and desires. (c) *Emotions and decision-making*: there are more and more findings showing an important correlation between decision-making under risk and emotional arousal (Mellers et al. 1997, Lopes 1987, Rottenstreich and Hsee 2001). These results are difficult to square with the belief–desire model that is silent on any emotional effects. (d) *The neural correlates of decision-making*: there is a tremendous amount of new findings about what parts of the brain are involved in decision-making (Lee 2006, McCoy and Platt 2005, Trepel et al. 2005, Tom et al. 2007, Platt and Glimcher 1999). These findings make it very implausible that decision-making is merely a matter of comparing beliefs and desires: it seems that it is a highly emotional affair.

Morsanyi and Handley 2012 and Nanay forthcoming b). But the pervasiveness of these influences jeopardizes the belief–desire framework as the general account of decision-making.

It would be tempting to add various extra elements to the existing belief–desire framework: maybe the general structure of decision-making is the belief–desire structure, but it's not the end of the story. Maybe there are some further (maybe affective) mental states that would also need to be postulated. So the suggestion would be that while we do make decisions by comparing the satisfaction conditions of our desires given various background beliefs, this process is further complicated by some other factors—maybe a dirtiness-detector state that influences the decision some way, or a teddy bear-detector state that influences the decision some other way. While we could of course add these further postulates thereby saving the general scientific research program of the belief–desire picture, it is difficult not to notice that these somewhat ad-hoc postulates—the only purpose of which is to save the research program from objections—are exactly what Lakatos identified as the mark of a degenerating research program.

Lakatos distinguished between progressive and degenerative scientific research programs (Lakatos 1970, 1974). A progressive research program does not contradict any new data, and makes new predictions and new explanations. A degenerative one sometimes does contradict new data, and makes no new predictions and new explanations. If a degenerative research program contradicts new data, this does not falsify the research program: there are many ways of modifying the research program in such a way that the contradiction disappears. These modifications, however, involve adding extra, ad-hoc, assumptions to the "core" of the research program, that serve only one purpose: to explain away the contradiction. These extra assumptions constitute the "protective belt" of a degenerative research program. The thicker the protective belt, the more likely it is that a research program is degenerative. The more new predictions and explanations a research program provides, the more likely it is that it is progressive (Nanay 2010g, 2011d). Lakatos argues that it is often worth being loyal to a degenerative research program for some time (as it may manage to recover), but if there is an alternative, progressive research program on the horizon, the rational thing to do is to jump ship and join the progressive one.

My claim is that the belief–desire model of decision-making shows all the typical signs of a degenerative scientific research program—it is mainly concerned with explaining away counterexamples rather than making new predictions, and it has accumulated a significant "protective belt" of more or less ad-hoc assumptions, the sole purpose of which is to explain away counterexamples. So we could continue pursuing this degenerative scientific research program, or else we could look for an alternative, progressive scientific research program. My preference would be the latter.

A further worry about my conclusion: maybe what changes in these examples is not our background beliefs and desires, but the weight that we assign to them. The suggestion would be the following: the warm coffee in my hand (or the dirty desk in front of me) does not change my beliefs and desires, but rather it makes me attend to certain beliefs or desires of mine and makes me ignore some others.

The first thing to note is that beliefs and desires are very different from the point of view of this argument. As beliefs are supposed to form a (by and large, most of the time) coherent network, shifting one's attention from one belief to the other should not result in any significant change in the outcome of the decision. In short, it is unlikely that focusing on one belief, rather than another, would be able to bring about the differences in decision-making that are demonstrated by the experiments I mentioned above.

Let us take the teddy bear experiment as an example. We have seen that holding a teddy bear (as opposed to merely seeing one) influences one's decisions about social behavior (Tai et al. 2011). How can the belief–desire model explain this? None of my beliefs change just because I hold the teddy bear instead of merely looking at it. But the proposal was that the subjects shift their attention from one belief (or beliefs) to another (or to other ones). How would this work? The proponent of the belief–desire model would need to postulate a (maybe unconscious) belief that is more salient in one of the scenarios than in the other—call this belief B*. Maybe B* is the belief that the teddy bear is warm and fuzzy. It could be argued that I have B* in both scenarios, but it is more salient in the one where I hold the teddy.

Now, in order for this postulation to save the belief–desire model, two claims need to be argued for. First, that B* is more salient when I hold the teddy than it is when I am merely looking at it. This may seem easy, but it

is not—there are many examples of differences in decision-making without any accompanying difference in phenomenology (see Newell and Shanks forthcoming for a summary of the literature on this). And if the change in salience/focus is to be cashed out in a way that does not appeal to phenomenology, then it is difficult to see what justification there is for postulating this difference—besides salvaging the belief–desire model, of course.

But the real problem is with the second claim that would need to be argued for: namely, that a difference in the salience of B* would make a difference in the subjects' judgment about something that has nothing to do with the contents of B*. Why would focusing on the belief that this teddy is fuzzy influence my judgment about others' (not teddy bears') social behavior? The belief–desire model could postulate even further mental states that would bridge the gap between B* and the decision-making, but this would look more and more like a textbook illustration of what degenerating scientific research programs do when they encounter an objection.

But maybe the difference between the two scenarios lies in the desires one has or attends to. And desires are different from beliefs: they are not supposed to form a coherent network—we can have (and we do have) contradicting desires. But this will not help the defenders of the belief–desire model of decision-making either, as in the case of many (in fact, most) of the experiments I mentioned here it seems that there is only one relevant desire (or pro-attitude) that would play a role in the decision-making process: in the coffee cup case, the only desire (or pro-attitude) that plays a role in the subjects' decision-making seems to be to answer the experimenter's question. The same goes for the pizza box study, and the studies of framing and order effects. But if these effects are manifest even if there is only one relevant desire (or pro-attitude), then the influences demonstrated by these studies clearly cannot be explained away with reference to shifting one's focus from one desire to another.

In short, the belief–desire model of decision-making has trouble accommodating a large number of empirical findings about how we actually make decisions. The main aim of this section is to outline an alternative account of decision-making.

A final and more general worry: couldn't we say that while the belief–desire model is a normative model of how we should make decisions, the reality is that we do not always make decisions the way

we should? In this case, all the experiments I mentioned in this section would count as deviations from the norm. A version of this claim is to say that the belief–desire model describes rational decision-making, but we are not always, and not fully, rational beings. Yet another version of this claim would be to say that while the belief–desire model is a good approximation of compensatory decision-making, our actual decision-making is often non-compensatory. And yet another version of this general claim is that while the belief–desire model is the right description of System 2, the automatic and mainly unconscious System 1 often overrides System 2 (Sloman 1996). The general moral of all of these distinctions (normative/descriptive, rational/irrational, compensatory/non-compensatory, System 2/System 1) is that the empirical findings I mentioned above do not challenge the belief–desire model, they only explain why it sometimes fails to apply to actual decision-making.

The problem with this objection is twofold. First, in order for us to be able to consider the actual cases to be deviations from a norm, the actual cases and the idealized/rational case shouldn't differ systematically and radically. But, arguably, this is exactly what the empirical results outlined above show. If the normative model of how we should make decisions is not even close to how we do in fact make decisions, then too much work is done by the biases and deviations from the norm—it becomes unclear whether the reference to the normative claims about how we should make decisions helps at all.

Second, while I do want to acknowledge the distinctions mentioned in the previous paragraph, and I would also want to acknowledge that some versions of the belief–desire model are better suited to explain decision-making than others, my problem is not with one version of the belief–desire model or another, but with the belief–desire framework as a scientific research program in general. My aim here is to look for alternatives to the scientific research program in general, and my reason for this is that, as we have seen, it is what Lakatos would call a degenerative scientific research program.

As Lakatos rightly observes, scientific research programs do not get falsified—and this is true of the scientific research program that explains decision-making in terms of beliefs and desires. They are abandoned if a better—that is, more progressive—scientific research program is being offered. My aim is exactly this: to offer a more

progressive scientific research program—one that aims to explain decision-making in terms of imagination.

4.4.2 Imagination and decision-making

My main aim is to outline an alternative model of decision-making that is consistent with the empirical findings presented in the last section. This alternative model emphasizes the role imagination plays in our decisions. To put it very simply, when we decide between two possible actions, we imagine ourselves in the situation that we imagine to be the outcome of these two actions and then compare these two imaginings.

Here is an example. You need to decide between two academic jobs: one of them is at a prestigious university in a not very nice small town, and the other is at a not very prestigious university in a great city. How do you decide? The belief–desire model would suggest that you have some desires (or other pro-attitudes) about how you want to live the rest of your life, and you also have some background beliefs; deciding between the two jobs is a matter of comparing the satisfaction of these desires given the background beliefs in the case of the two choices. We have seen that this account faces various empirical problems.

My suggestion is much simpler. When you decide between the two jobs, you imagine yourself in the situation that you imagine to be the outcome of your decision one way or the other. You imagine yourself at the prestigious university surrounded by great colleagues and doing excellent research in a sleepy small town, spending the evenings working or with colleagues. You also imagine yourself at the not so prestigious university, spending every night out in cool restaurants and at various cultural events, to return to teaching the next day among your mediocre colleagues and not so bright students. Then you compare these two imaginative episodes, and the one you prefer will be the course of action to follow.

A somewhat similar model was put forward in a series of papers by Jonathan Evans (see especially Evans 2007). He argues that:

> We need somehow to imagine the world (in relevant respects) as it might be following a particular choice or action under our control and decide how much we would like to be living in it. Moreover, we need to conduct a set of thought experiments for each possible action and compare their evaluations.
>
> (Evans 2007, p. 12)

While this may sound similar to the proposal I was making, there are also important differences. A salient difference is that Evans's main concern

is about how we *should* make decisions—thus the emphasis on how we should conduct these thought experiences for each possible action and compare their evaluations. I am interested in how we do in fact make decisions.

More importantly, as he makes clear elsewhere (especially in Evans and Over 2004, p. 9), Evans believes that this decision-making process by means of hypothetical thinking is what happens in System 2, and it would be a rational and reliable way of making decisions this way if System 1 were not giving rise to all kinds of biases. The method Evans describes is rational and reliable. The decision-making process I talk about is neither.

I need to say more about just what imaginative episodes are involved in our decision-making. First, note that the concept of imagination plays (at least) two roles here: you *imagine* yourself in a situation you *imagine* to be the outcome of your decision. Let us start with the first of these: imagining oneself in a future situation. This kind of imaginative episode is often described as "imagining from the inside."

But what does imagining another person, X, from the inside mean? There are two different views about what this concept entails:

(1) Imagining having X's experiences: X occurs in the content of my imaginative episode, I myself may not.
(2) Imagining being in X's situation: the imaginative episode is all about myself; X herself does not occur in the content of my imaginative episode.

(1) has been the dominant view of "imagining from the inside" (Currie 1995b, p. 153, Neill 1996). An influential version of (1) is Kendall Walton's account of imagining from the inside: when I imagine X from the inside, I imagine experiencing what I think X experiences (Walton 1990, pp. 255, 344).

Arguably, Adam Smith was the first philosopher who held a version of (2) (Smith 1759/1976, see also Nanay 2010e. A detailed contemporary formulation of (2) is in Gaut 1999). The crucial difference from (1) is that X is not part of the content of this imaginative episode. Only I myself and X's situation are.[11]

---

[11] Not all accounts of imagining from the inside fall clearly into one of the two categories I differentiated above. Gregory Currie's account, for example, is ambiguous between (1) and (2). Currie does not talk about imagining from the inside, but about what he calls

The imaginative episode that plays a role in our decision-making is to be characterized as (2). But X is not another person; it is your future self. You imagine yourself in the situation you imagine you will be in, given the decision you are making. This is an instance of self-imagining.

Further, to make things more complicated, while you know some things about the situation you take your decision to lead to (that is, living in a small college town, etc.), you have to rely on your imagination yet again: there are plenty of details in this picture of living in a small college town that need to be filled in, and where your knowledge cannot help any more (what you'll do on Saturday nights, what your house will be like, whether this or that potential colleague is sociable)—these details need to be filled in by your imagination. So you *imagine* yourself in a situation that you *imagine* to be the outcome of your decision.

But there is yet another important complication that may force us to allude to imagination not twice, but three times in this model. Imagining oneself in X's situation is itself ambiguous. It can be interpreted either as imagining being in X's situation, or as imagining being X in X's situation. This distinction is not new (Williams 1973, Wollheim 1973, 1974, Reynolds 1989, Velleman 1996), and it is not as straightforward as it may seem (see, for example, Gordon 1995a, 1995b, Recanati ms).

It is crucial to be clear about this distinction for the purposes of my proposal about decision-making, as it matters a great deal whether I imagine myself—that is, my (present) self—in a future situation that I take to be the outcome of my decision, or whether I imagine my future self (or what I take to be my future self) in that situation.

Who do you imagine when you imagine yourself living in a small college town? Do you imagine your future self or your present self in that situation? It seems that neither imaginative episode would give us particularly reliable information (on the basis of which we could make the "right" decision), but for very different reasons. In order to make the

"secondary imagining." Sometimes he characterizes secondary imagining in ways that correspond to (1): as imagining "the experience of the character" or "in imagination, feeling what the character feels" (Currie 1995b, p. 153). Sometimes, however, he gives formulations of secondary imagining that are very similar to (2): "as a result of putting myself, in imagination, in the character's position, I come to have imaginary versions of the thoughts, feelings and attitudes I would have in that situation" (Currie 1995b, p. 153). As these two different conceptions of secondary imagination occur on the very same page, we have good reason to suppose that he takes "X's situation" to be so broad as to make (2) collapse into (1). See also Currie 1995a, and Currie and Ravenscroft 2002.

"right" decision, you would need to imagine your future self in that situation—it is irrelevant how your present self would feel in that situation, as your present self will never be in it.

But in order to imagine your future self in that situation, you would need to have some idea about how your future self may be different from your present self—not a straightforward thing to do. In fact, while you can imagine what your future self will be like, you have no reliable information about what it will be like.[12] It seems that you need to rely on imagination at this point yet again. If, however, you imagine your present self in a future situation, then it is not clear that this imaginative episode has any real bearing on how your future self would indeed feel in this future situation. Again, this imaginative episode will not give you reliable information on which the optimal decision could be based.[13]

To sum up, there are three points where imagination plays a role in the decision-making process. You imagine what you imagine to be your future self, being in a situation that you imagine to be the outcome of your decision. As none of these three episodes of imagination can be considered reliable, decision-making is extremely unlikely to yield the optimal outcome reliably.

Of course, we can, and often do, make the optimal decision, especially when it is about some decision-problem that concerns the near future, or one we encounter often—for example, where to go to get a decent cup of coffee in the neighborhood. But we have no reason to think that our decisions in general, and especially decisions that really matter to us, have much of a chance at yielding the optimal outcome.

Note the contrast with the belief–desire model: the scientific research program that aims to explain decision-making in terms of beliefs and desires. The belief–desire model takes rational and reliable decision-making to be the paradigmatic case, and it aims to explain deviations from this rational norm by appealing to various biases. The picture I want to replace this scientific research program with proceeds in the opposite direction. It takes our actual, very unreliable, not at all

---

[12] Further, we have some (rather depressing) empirical evidence that people systematically ignore the possibility that their future self could be different from their present self (see Quoidbach et al. 2013).

[13] See also the vast literature on affective forecasting on this topic.

rational ways of making decisions to be paradigmatic.[14] Its starting point is not the way we should make decisions, but the way we do in fact make decisions. And, as it turns out, the way we do in fact make decisions is extremely unreliable.

The main problem with the belief–desire account of deliberative actions was that it cannot accommodate a number of important recent empirical findings about what influences our actual decision-making. In conclusion, I need to show that the same empirical findings about decision-making are better explained by the account I proposed above.

Whether you hold a cup of hot coffee in your hand or whether you are surrounded by dirty pizza boxes is unlikely to change your background beliefs and desires. Can they influence your imaginative episodes? The answer is that they definitely can: it has been known for a long time that imagination is extremely sensitive to all kinds of external and seemingly irrelevant influences (see the *locus classicus*, Feingold 1915, and see also more specific findings about the influences on mental imagery, Segal 1972, Segal and Nathan 1964, Nanay 2010c, Raftopoulos forthcoming, about which I will say more in the next chapter). In other words, if, as I suggested, our decision-making is a matter of imagining ourselves in various situations, what we should expect is that decision-making is as sensitive to external, seemingly unrelated influences as imagination itself is. And, as the recent empirical findings show, this is indeed the case.

Much of the rhetoric of this section was based on painting the belief–desire model as a degenerating research program. But why should we think of the imagination-centered research program as progressive? A progressive research program explains novel phenomena, or novel features of known phenomena. Does my imagination-centered research program do that? I would argue that it does. It can give a nice and simple explanation for some of the differences between the way we make decisions and the way it would be rational to make decisions.

One widely reported bias in decision-making is that we systematically overestimate the probability of certain events and underestimate the

---

[14] In this sense, Evans's work, which in many respects would constitute an instance of the scientific research program that I propose to replace the belief–desire model with (as it explains decision-making in terms of imagination and not of beliefs and desires), does resemble the belief–desire model more than mine: Evans's starting point is, as we have seen, also the way we should make decisions (see especially Evans and Over 2004, p. 12).

probability of other events. One famous example is about people's estimation of what causes the most deaths, say, in the US. Most people widely overestimate the number of deaths resulting from homicide or fireworks, and underestimate the number of deaths caused by asthma or diabetes (Slovic et al. 1976). The popular science explanation for this is that we read about and watch news footage of firework disasters all the time, but are much less exposed to news items about people dying of asthma. But this explanation needs to be supplemented by a view about the ways in which the media exposure influences our decision-making.

A very straightforward view would be to say that we find it easier to imagine those scenarios that we have encountered (or read about) more often than those that we have never encountered (and rarely read about). So the difference is explained in terms of what is more imaginable. The same strategy then could be used to explain a wide variety of biases. Just one quick example: most people want a higher payoff in a lottery where there is one winning ticket out of 100 tickets if they know that all the other 99 were bought by one and the same person, call her Jennifer (Walker 1992, Wagenaar 1988). Why is that? Because it is easier to imagine that Jennifer has the winning ticket. Again, the difference consists of what one can imagine easily. This is a feature of our biases in decision-making for which the imagination-centered scientific research program can provide a much simpler explanation than the belief–desire model—which is a sign of a progressive research program.

A final question: just how radical a change am I proposing here? I would say not a very radical one at all. The concept of imagination has been used very consistently in the standard belief–desire model literature on decision-making. Just one example: the concept of "availability bias" (Tversky and Kahneman and Tversky 1973, Tversky and Kahneman 1982) really amounts to bias in imaginability. Tversky and Kahneman's original formulation is that "there are situations in which people assess the [. . .] probability of an event by the ease with which instances or occurrences can be brought to mind" (1982, p. 9). It is difficult to see what could be meant by "being brought to mind" if not imagination. My suggestion merely amounts to making this concept of imagination play a more central role in accounting for decision-making.

To summarize, we have good reasons to replace the belief–desire model of decision-making with a model where imagination plays a crucial role. And this would push the belief–desire model even further away from the

monopoly it once held in thinking about the mind. We have seen that it is unlikely that the belief–desire model is the most fruitful way of describing the mental processes that lead to our simple actions like brushing our teeth or tying our shoelaces. But, if the argument I presented in this section is correct, then it may not be a particularly fruitful way of thinking about the mental processes that lead to even our most sophisticated and deliberated actions, like deciding between two job offers.

I would like to close with a literary quote that could be considered to give an account very similar to the one I argued for here—from Marcel Proust: "Since my parents had told me that, for my first visit to the theatre, I should have to choose between these two pieces, I would study exhaustively and in turn the title of one and the title of the other (for those were all that I knew of either), attempting to snatch from each a foretaste of the pleasure with that latent in the other title."[15]

---

[15] Proust, M. (1928) *Swann's Way* (trans. C. K. Scott Moncrieff). New York: Modern Library, p. 102.

5

Pragmatic Mental Imagery

> We need to confront vague ideas with clear pictures.
>
> Jean-Luc Godard, *La Chinoise*, 1967.

5.1 Mental imagery

In Chapters 3 and 4, I outlined some important consequences of pragmatic representations for philosophy of perception and for philosophy of action. In Chapters 5 and 6, I consider some important variants of pragmatic representations: pragmatic mental imagery and vicarious perception.

Pragmatic representations attribute action-properties to objects: that is, objects in front of us, perceptually. I want to consider a kind of mental state that, in some ways, is very similar to pragmatic representations: mental states that attribute action-properties to objects by means of mental imagery. I call mental states of this kind *pragmatic mental imagery*.

I will first characterize mental imagery and analyze the ways in which it could be thought to be similar to perceptual states, and then identify a thus far underexplored kind of mental imagery: pragmatic mental imagery. I then argue that mental states of this kind play an important role in our mental life.

Here is a relatively general characterization of mental imagery:

Mental imagery refers to all those quasi-sensory or quasi-perceptual experiences [. . .] which exist for us in the absence of those stimulus conditions that are known to produce their genuine sensory or perceptual counterparts, and which may be expected to have different consequences from their sensory or perceptual counterparts. (Richardson 1969, pp. 2–3)

It is easier to explain mental imagery in the visual sense modality, which is the one I will mainly be focusing on. A paradigmatic case of visual imagery would be closing one's eyes and imagining seeing an apple "in the mind's eye" (see Kosslyn 1980, Kosslyn et al. 1995, 2006, see also Ryle 1949, Section 8.6, Currie and Ravenscroft 2002). The equivalent of visual imagery in other sense modalities would be auditory, tactile, or olfactory imagery. I will use the term "mental imagery" to refer to all of these.

It is important to point out that visual imagery does not necessarily imply visualizing. We usually think of visualizing as an active, intended act. Having mental imagery, on the other hand, can be passive and is not necessarily intended. Visualizing is one way of having mental imagery, but it is not the only way. We can have mental imagery even if we are not trying to visualize anything—when, for example, we are having involuntary flashbacks to some scene that we have seen earlier. This is especially clear if we shift our attention to the auditory sense modality and consider earworms: tunes that pop into our heads and that we keep on having auditory imagery of, even though we do not want to.

Mental imagery then is a mental state that attributes, quasi-perceptually, properties to a scene or object that this scene or object is not perceived as having—where "quasi-perceptually" just means in a way that perceptual states also attribute properties.[1]

This can happen in two importantly different ways. Mental imagery can attribute, quasi-perceptually, properties to an imagined object. This is what happens if I close my eyes and visualize an apple. But mental imagery can also attribute, quasi-perceptually, properties to a perceived object—properties that we do not perceive this object as having. This distinction will play an important role below.

Although the stereotypical example of mental imagery is closing one's eyes and visualizing an apple, as we have seen, this is not the only kind of mental imagery. Mental imagery can be involuntary. Dreams and hallucinations are involuntary, yet they are also normally interpreted as examples of mental imagery. When we have a dream, we attribute properties

---

[1] I said in Chapter 3 that much of what I say in this book is neutral between the representationalist and the relationalist ways of thinking about perception. I also said that this is not true of everything I say in this book. The present chapter is the exception—I will make heavy use of the representationalist framework here when I am characterizing mental imagery and its relation to perception (especially in the next section). So this is a warning to the relationalists among the readers: they may want to just skip this chapter.

quasi-perceptually to a scene that is not in fact there in front of us. The same goes for hallucination.

If the argument I presented in Nanay (2010c) is correct, then amodal perception (or at least most instances of amodal perception, see Briscoe 2011)—that is, the representation of unseen parts of objects we see—is also a sub-category of mental imagery, where, again, we attribute properties quasi-perceptually to a part of the perceived object that is not visible. Further, mental imagery is also taken to be a necessary feature of episodic memory, as it has been pointed out that the loss of the capacity to form mental imagery results in the loss (or loss of scope) of episodic memory (Byrne et al. 2007, see also Berryhill et al. 2007's overview).

Having some kind of mental imagery of an apple should be differentiated from imagining that there is an apple in the kitchen, an imagining episode often labeled as propositional imagination. The latter is a propositional attitude, whereas the former is a quasi-perceptual process (see Nanay 2009a, 2010c, Van Leeuwen 2011).

The connection between mental imagery and propositional imagination is not at all clear. Is mental imagery necessary for propositional imagination? Is it sufficient? The sufficiency claim can be dismissed as we have mental imagery when we are dreaming, but dreaming is very different from propositional imagination. Episodic memory also seems to presuppose the exercise of mental imagery, but it is, again, very different from propositional imagination.

The necessity claim is more complicated (see, for example, Kind 2001, Chalmers 2002, Baron-Cohen 2007, Byrne 2007). A tempting proposal is that the content of propositional imagination is built up or derived from mental imagery (Baron-Cohen 2007, p. 104): to imagine that mermaids have blue eyes, for example, requires some transformation of the visual imagery of a fish and a woman. But if this is so, then mental imagery is necessary for propositional imagination (moreover, this may constitute the difference between propositional imagination and supposition: only the former requires mental imagery). But not everyone agrees that mental imagery is necessary for propositional imagination (see Chalmers 2002, p. 151, for a detailed objection). I do not want to take sides in this debate.

A couple more clarifications about mental imagery are worth adding. First, mental imagery can be conscious or unconscious (in this respect, I depart from the Richardson quote above, which could be interpreted as

stating that mental imagery is necessarily conscious). Given the various similarities between perception and mental imagery (see Section 5.2), if perception can be conscious or unconscious, then it is difficult to see what would prevent one from having both conscious and unconscious mental imagery. I argued in Nanay (2010c) that we represent unseen parts of perceived objects by means of mental imagery, and we normally do so unconsciously. I will not rely on this argument in this book, but if it is true then we must have unconscious mental imagery most of the time.

Further, although when we close our eyes and visualize an apple, we tend to visualize this apple as being located in some kind of abstract space, this is not a necessary feature of mental imagery. We can have mental imagery of something as located in our own space (Martin 2002a, p. 410, Nanay 2010c). In other words, we can look at an empty chair in our egocentric space and have mental imagery of our friend sitting in it. Further, mental imagery sometimes attributes a property to a perceived object in our egocentric space—for example, when I look at a brown chair and imagine it to be red. In this case, my mental imagery attributes, quasi-perceptually, a property to the chair that this chair is not perceived as having—just as the working definition of mental imagery suggests.

5.2 Mental imagery versus perception

Perception and mental imagery are very similar in a number of respects. There is, for example, an almost complete overlap between the brain regions involved in perception and the brain regions involved in mental imagery, which suggests that the mental processes that make perception possible are the very same mental processes that make mental imagery possible (see, for example, Kosslyn et al. 2006, Shepard and Metzler 1971). Further, the patterns of cortical activation are also similar in perception and mental imagery (Page et al. 2011). Note that this claim is not contested in the grand "mental imagery debate" (see Tye 1991): even those who argue that mental imagery is propositional (for example, Pylyshyn 1999) would grant the neurological similarities between the two processes.

Further, conscious perception and conscious mental imagery have very similar phenomenal characters. If I visualize a red apple, the phenomenal character of my experience will be very similar to that of my

experience if I see one. The phenomenal similarity between seeing and visualizing seems intuitively salient, but for those (like myself) who mistrust intuitive evidence, there is empirical evidence for this similarity. In the Perky experiments, subjects looking at a white wall were asked to visualize objects while keeping their eyes open. Unbeknownst to them, barely visible images of the visualized objects were projected on the wall. The surprising finding is that the subjects took themselves to be visualizing the objects—while in fact they were perceiving them (Perky 1910, Segal 1972, Segal and Nathan 1964). The standard interpretation of this experiment is that if perceiving and visualizing could be confused under these circumstances, then they must be phenomenally very similar (but see Hopkins 2012's criticism, and Nanay 2012b's response).

How can we explain this phenomenal similarity between perception and mental imagery? One (not the only one—see the literature on the "dependency thesis," Martin 2002a, Noordhof 2002, Smith 2006, pp. 53–4, Peacocke 1985) straightforward way of explaining this phenomenal similarity is to argue that the content of perceptual states and of mental imagery is similar. If the content of these two different kinds of mental states are similar, it should not come as a surprise that their phenomenal character is also similar. Following Noordhof (2002, p. 439), I call this explanatory scheme the *similar content view*.

The plausibility of this explanatory scheme clearly depends on the way we think about perceptual content. Perceptual content is a semi-technical term, and depending on how we conceive of perceptual content, we end up with very different versions of the similar content view. This means that the similar content view needs to be supplemented with an account of the nature of perceptual content and of the content of mental imagery.

Here is one way of thinking about perceptual content that does not seem to be particularly promising when it comes to fleshing out the similar content view. If we equate perceptual content with the object the perceptual state is about, then the similar content view will amount to saying that visualizing a green chair and seeing one have similar phenomenal characters, because the green chair that I see is similar to the green chair that I visualize.

The problem is that this view does not seem to have the resources to explain in what sense the two kinds of contents are similar. The green chair that I visualize may not exist, whereas the one I see does. In short,

under this conception of perceptual content, the contents of perceptual states and of mental imagery are very different indeed: an actual token object versus a potentially nonexistent object. One way of salvaging this way of thinking about perceptual content would be to say that the image of the imagined chair is similar to the image I see when I am looking at the chair. But this view was often ridiculed (most influentially by Gilbert Ryle, in Ryle 1949) as the "special status pictures" view of imagination.

It is important to note that the similar content view does not need to postulate the existence of such "special status pictures." Another more general version of the similar content view is that the properties we attribute in imagination to the visualized object are similar to the properties we attribute in perception to the object seen (Ishiguro 1967, cf. Kind 2001). Unlike the "special status picture" theory, this version of the similar content view is still quite popular (Currie 1995a, pp. 36–37, Currie and Ravenscroft 2002, p. 27, Noordhof 2002).

Here is my version of the similar content view. Consider the following very simple, and not particularly controversial, way of thinking about perceptual content (Nanay 2010a). Our perceptual apparatus attributes various properties to various parts of the perceived scene, where I take the perceived scene to be spatially (and not propositionally/syntactically) organized, in the way Peacocke's scenario content is (Peacocke 1986, 1989, 1992, Burge 2010). Perceptual content is constituted by the properties that are perceptually attributed to the perceived scene.

In order to maintain the generality of this account of perceptual content, I will say nothing about whether these properties are tropes or universals (Nanay 2012d), or whether this content is structured in a Fregean or Russellian manner (Nanay forthcoming f), for example (see Chapter 3 for more on this). The question I want to explore here is what degree of determinacy these perceptually attributed properties have.

One way of characterizing the relation between properties is the determinable–determinate relation (Johnston 1921, Funkhouser 2006). To use a classic example, being red is determinate of being colored, but determinable of being scarlet. There are many ways of being red, and being scarlet is one of these: for something to be scarlet is for it to be red in a specific way. If something is red, it also has to be of a certain specific shade of red: there is no such thing as being red *simpliciter*.

The determinable–determinate relation is a relative one: the same property—for example, of being red—can be the determinate of the

determinable being colored, but the determinable of the determinate being scarlet. Thus, the determinable–determinate relation gives us a hierarchical ordering of properties in a given property-space. Properties with no further determinates, if there are any, are known as super-determinates.

Some of the properties we perceptually attribute to the perceived scene are determinates, or even super-determinates. Some others, on the other hand, are determinable properties. We know that our peripheral vision is only capable of attributing extremely determinable properties. But even some of the properties we perceptually attribute to the objects that are in our fovea can be determinable.

It has been argued that if we accept this way of thinking about content, then perceptual attention is a necessary feature of perceptual content (Nanay 2010a, Jagnow 2011, Nanay 2011e). More precisely, attention makes (or attempts to make) the attended property more determinate (see also Yeshurun and Carrasco 1998 for empirical evidence, and Stazicker 2011 for a philosophical summary).[2] If I am attending to the color of my office telephone, I attribute very determinate (arguably super-determinate) properties to it. If, as is more often the case, I am not attending to the color of my office telephone, I attribute only deter-minable properties to it (of, say, being light-colored or maybe just being colored). In short, attention makes (or attempts to make) the perceived property more determinate.

An important clarification: a shift of visual attention is not to be confused with eye movement. It is possible to shift one's visual attention without any accompanying eye movement—this is a widely researched phenomenon of the "covert shift of attention" (Posner 1980, 1984, Posner et al. 1984, see also Findlay and Gilchrist 2003). But more often the shift of attention is accompanied by eye movement, which, following the literature, I call an "overt shift of attention." Both in the case of overt and of covert shifts of attention, the determinacy of the attended property changes.

This, I believe, is a simple, and not particularly controversial, account of perceptual content. But what is the content of mental imagery?

---

[2] It may be worth noting that this way of thinking about the relation between attention and perceptual content provides a direct explanation for the inattentional blindness phenomenon (see Mack and Rock 1998, Simmons and Chabris 1999).

My answer is that the content of mental imagery is exactly the same as the content of perceptual states.

More precisely, our imagery attributes various properties to various parts of the imagined scene. The content of imagery is the sum total of the properties attributed to the imagined scene. Some of these properties are determinates, or even super-determinates. Some others are determinables. Attention makes (or tries to make) the attended property more determinate.

What, then, is the difference between perceptual content and the content of mental imagery? The only difference concerns where the extra determinacy comes from. As we have seen, both in the case of perceptual content and in the case of mental imagery, attention makes the attended property more determinate. This increase in determinacy in the case of perception comes from the sensory stimulation: if I am attending to the color of the curtain in the top left window of the building in front of me, this color will be more determinate than it was when I was not attending to it. This difference in determinacy is provided by the world itself—I can just look: the exact shade of the curtain's color is there in front of me to be seen.

In the case of mental imagery, this difference in determinacy, in contrast, is not provided by sensory stimulation, for the simple reason that there is no sensory stimulation that would correspond to what I visualize: if I visualize the house I grew up in and you ask me to tell what exact color the curtain in the top left window was, I can shift my attention to that color and I can even visualize the exact color of the curtain. However, this increase in determinacy is not provided by the sensory stimulation (as I don't have any), but by my memories (or what I take to be my memories) or my beliefs/expectations.

Let's consider another example where the increase in determinacy is not provided by my memories (or by what I take to be my memories), but by my expectations. Suppose that I order a steak in a restaurant, and I have mental imagery of the meal the waiter is about to bring me. I can shift my attention around here as well—I can attend to the texture of the meat, for example. This, again, would entail an increase in the determinacy of this texture, but this increase is not provided by my memories, but by my expectations—in this case, expectations based on my belief that I asked the steak to be cooked medium-rare.

Clarifications: first, my account is not committed to there being a clear-cut distinction between perception and mental imagery. In the modified Perky experiments (Segal 1972), the picture projected on the wall and the image the subjects were asked to visualize were different, resulting in an interesting juxtaposition of the two images. In this case, it would be difficult to tell whether the subject perceives or exercises mental imagery—she does both (see Trehub 1991 for some further experiments involving mixed perception/mental imagery). The fact that according to my account the structure of the content of these two mental episodes is the same makes it easy to account for mixed cases like this (other somewhat different examples of mixed perception/mental imagery are given in Martin 2002a, p. 410). The increase in determinacy is provided by both the sensory stimulation and our memories/beliefs in these cases.

Second, my claim is not that attention makes the attended property more determinate, but that it makes *or tries to make* the attended property more determinate. It does not always succeed. And this is so both in the case of perceiving and in the case of visualizing. When I attend to something that I see in the periphery of my visual field and I cannot move my eyes, the shift of my attention tries to make the properties of this object more determinate, but because this object is, and continues to be, in the periphery of my visual field, I will not succeed. The same goes for mental imagery. If I am asked to visualize my first credit card and attend to its color, I may just simply not remember, and in this case, although attention tries to make the attributed property more determinate, it may not succeed.

In short, the difference between perceptual content and the content of mental imagery is not a difference between the structure of these contents—they have the very same structure. The difference is between the dynamics of how the represented properties, and, importantly, the determinacy of the represented properties, change in response to the allocation of attention. The difference is not between what perceptual content and the content of mental imagery are, but between the ways they change.

It is important to emphasize that the claim is not that the properties attributed in the content of mental imagery are less determinate than the ones that are attributed in perceptual content. The properties that constitute the content of mental imagery can be very determinate indeed (and most of the properties that constitute perceptual content are not

particularly determinate—see Dennett 1996). The claim is that the difference between the content of these two mental states is the way this determinacy comes about.

I argue elsewhere (Nanay ms) that this way of cashing out the similar content view is more explanatorily powerful than either the dependency thesis or other versions of the similar content view. But this is not the aim of this chapter. The aim is to show that the claim that mental imagery attributes properties quasi-perceptually can be substantiated—if the reader is not convinced by my specific way of substantiating this claim, she can plug in any other account of the content of mental imagery so long as mental imagery is understood as attributing properties quasi-perceptually.

5.3 Pragmatic mental imagery

But if mental imagery attributes properties quasi-perceptually, then we can consider a special subcategory of mental imagery where these attributed properties are action-properties. I call this subcategory of mental imagery that attributes action-properties quasi-perceptually "pragmatic mental imagery."

Here is an example. Suppose that you are looking at a cup in front of you and then you reach out to lift it up. As we have seen, in order for you to be able to do this, you must represent the cup as having some action-properties: some properties the representation of which is necessary for the performance of this action. You need to represent the cup as having a certain shape, otherwise you could not approach it with the appropriate grip size. You need to represent it as having a certain weight, otherwise you could not exert the appropriate force. And so on. In short, in order to lift the cup, you need to attribute these action-properties to it perceptually. The representational component of the immediate mental antecedent of the action of lifting the cup is the pragmatic representation that attributes these action-properties.

Now suppose that you are looking at the same cup. You want to lift it, but before you do so you close your eyes for ten seconds and then reach out to lift it with your eyes closed. What is the immediate mental antecedent of this action? One tempting suggestion would be to say that the immediate mental antecedent is the perceptual state you had

before you closed your eyes. It is this perceptual state that guides your action with a ten-second delay. This may in fact be so in some cases.

But not always. Suppose that after you close your eyes, I tell you that I have moved the cup towards you by ten centimeters. And then you reach out with your eyes closed and pick up the cup. What is the immediate mental antecedent of your action now? It is difficult to see what else it could be but the mental imagery you form on the basis of your earlier perceptual state, as well as the verbal information I have given you. But in order for this mental imagery to allow you to perform the action, it needs to attribute action-properties to the cup that will allow you to reach in the right direction, with the right grip size, and with the right force. The immediate mental antecedent of this action is your pragmatic mental imagery.

Let us see what kind of mental imagery this pragmatic mental imagery is. We have seen that mental imagery can be voluntary or involuntary, conscious or unconscious, and that it may or may not localize the imagined object in one's egocentric space. How about pragmatic mental imagery? In the example I have just given, it is likely to be conscious and voluntary: you need to make an effort and consciously visualize the cup in the location where I told you I have shifted it. But there may be other cases of actions guided by pragmatic mental imagery where this pragmatic mental imagery is neither voluntary nor conscious. For example, when you turn on the light in a pitch-dark room that you know so well that you do not have to deliberately and consciously visualize the switch—you switch it on almost automatically. In order for you to be able to perform this action, you need to attribute action-properties to the light switch, and you are likely to do so by means of unconscious and involuntary mental imagery.

In short, like pragmatic representation, pragmatic mental imagery can also be conscious or unconscious. But, like pragmatic representation (and unlike some other forms of mental imagery), pragmatic mental imagery also needs to represent the object it attributes action-properties to in one's egocentric space. In other words, unlike the case of closing one's eyes and visualizing an apple in an "abstract" space, pragmatic mental imagery attributes action-properties to an object (perceived or unperceived) that is located in your egocentric space. Otherwise it could not attribute a spatial location property to this object that would allow you to perform actions with it.

Further, pragmatic mental imagery (like mental imagery in general) can attribute action-properties to an imagined (that is, not perceived) object, or it can attribute action-properties to a perceived object—action-properties that we do not perceive this object as having. Only the latter form of pragmatic mental imagery will be able to help us perform actual actions—attributing action-properties to imagined objects could help us perform actions with these imagined objects, but not with actual ones. But the attribution of action-properties to imagined objects is also a form of mental imagery, one that will play an important role in the next section.

It is important to point out that pragmatic mental imagery is not the same as motor imagery (Jeannerod 1997, Currie and Ravenscroft 1996). Pragmatic imagery represents the (imagined) object as having action-properties. Motor imagery, in contrast, represents how our body would move if we were to move. These two forms of imagery are directed at different things. Motor imagery is an important and interesting mental phenomenon, but not the one I will be focusing on in this chapter.

Further, just as the perceptual attribution of action-properties (that is, a pragmatic representation) can be accompanied by the awareness of thick action-properties, the attribution of action-properties in mental imagery (that is, pragmatic mental imagery) is also often accompanied by the awareness of thick action-properties. If I visualize a cake, then my pragmatic mental imagery may attribute action-properties to this imagined cake, but I may also experience the cake as having some thick action-properties (like being edible, for example).

In Chapter 3, I argued that whatever the exact implementation of pragmatic representations is, it is likely to include large chunks of the dorsal visual subsystem. Given that we also know that the neural circuits responsible for our perception are also responsible for our mental imagery (see Kosslyn et al. 2006's summary), putting these two claims together strongly suggests that the implementation of pragmatic mental imagery must include the dorsal visual subsystem. And, as it turns out, it does. The dorsal stream is often active during mental imagery—especially when the imagery is supposed to aid the subject in the performance of a task—for example, in the mental rotation task (see, for example, Mellet et al. 1998, Podzebenko et al. 2002).

We have seen that pragmatic representations are very old in evolutionary terms, and that this is not at all surprising given that natural

selection cares about the success of the actions we perform and that pragmatic representations are by definition geared towards helping us to perform actions successfully (see also Nanay 2013). We can make a parallel argument about pragmatic mental imagery.

It has been widely debated why mental imagery evolved at all. Was it because it helped us in "mental time travel"—remembering (Patten 1920, pp. 291–3) and planning (Mithen 2001, p. 33)? Was it because it helped us in pretense (Byrne 1998, especially p. 117 ff., Mithen 2001, pp. 33–4)? Or did it allow us to rehearse probable but non-actual situations offline (Tooby and Cosmides 2005, pp. 60–1)? Or facilitate our creativity (Nettle 2001, p. 141)? Many of these explanations are specific to humans, while we have evidence that mental imagery is quite widespread among vertebrates (Horridge et al. 1992, Nieder 2002, Nieder and Wagner 1999, Regolin and Vallortigara 1995). Even pigeons, for example, have mental imagery (Rilling and Neiworth 1987, Neiworth 1992, see also Oakley 1985 for a summary).

But attributing action-properties to momentarily invisible objects, or to those parts of perceived objects that are not visible to us at the moment, has obvious adaptive value.[3] Pragmatic representations help us to localize entities that are important for our survival in our egocentric space: food, predators, and potential mates. They are clearly selectively advantageous. Pragmatic mental imagery does the same thing: it helps us to localize *occluded* entities that are important for our survival in our egocentric space. If we are looking for the evolutionary origins of mental imagery, pragmatic mental imagery seems to be a very good starting point.

To sum up, pragmatic mental imagery can serve as the immediate mental antecedent of our actions as much as (but perhaps less efficiently

---

[3] If we accept, as I argued in Nanay 2010c, that we represent the occluded parts of perceived objects by means of mental imagery, then this becomes even clearer: it is a good idea to attribute action-properties to the animal I see only the tail of (whether it is predator or prey). Being able to localize the unseen parts of an animal hiding in a bush is an extremely survival-enhancing skill. As W. S. Ramachandran writes: "Why do illusory contours exist and how do they influence subsequent visual processing? In a world that is so rich in real contours it is hard to see what evolutionary advantage would accrue from the ability to construct illusory edges. But consider an arboreal primate trying to detect a leopard seen against the background of dense foliage. To this creature, segmenting the scene using illusory contours may be of vital importance as an anti-camouflage device. Many animals have developed elaborate splotchy markings in order to break their outlines [. . .], and the ability to perceive illusory contours may have evolved specifically to defeat this strategy" (Ramachandran 1987, p. 95).

than) pragmatic representations can. In Chapter 2, the main claim was that the (representational components of the) mental antecedents of our actions are pragmatic representations. I briefly alluded to the possibility of mental imagery playing the same role. In this chapter, I fleshed out this brief allusion. But then my claim in Chapter 2 needs to be amended slightly: the (representational component of the) mental antecedent of one's action is either pragmatic representation or pragmatic mental imagery.

But the aim of this chapter is not merely to clarify a potential ambiguity identified in Chapter 2. As it turns out, pragmatic mental imagery plays a very important role in our mental life: for example, in our pretense actions. This is the topic I now turn to.

5.4 Pretense

Pragmatic representations play a key role in bringing about actions. My aim in this section is to explore what role pragmatic mental imagery plays in bringing about pretense actions. When my daughter pretends to make pizza using her blanket and stuffed animals, this is a pretense action. One recently widely discussed question is what makes it a pretense action: what are its mental antecedents?

The modern *locus classicus* is Nichols and Stich (2003). They argue that pretense actions can be explained if we add a "possible world box" to the mental apparatus that brings about real (not pretend) actions. We imagine that x is F. This is what becomes the content of our "possible world box." This possible world box then gives rise to a conditional belief with contents like "if x were F, then I would do A." And this conditional belief, together with the desire to behave as if x were F, motivates the performance of the pretense action to do F. To use the pizza-making example, in the possible world box we have "I am making pizza," which gives rise to the conditional belief "if I were to make pizza, I would distribute the toppings evenly on the surface," which in turn, together with the desire to behave as if I were to make pizza, triggers the pretense action (of distributing stuffed animals evenly on the blanket).

This way of describing the mental underpinnings of pretense actions has been criticized for various reasons. According to the most important alternative model, the pro-attitude that leads to the pretense action is not

an actual desire, but an "imaginary desire": an i-desire (Doggett and Egan 2007), desire-like imagination (Currie and Ravenscroft 2002), or wish (Velleman 2000). And it is not the case that the imagination (the possible world box) triggers a real (but conditional) belief that in turn triggers the pretense action. The pretense action is triggered by imagination (belief-like imagination) directly. Thus, the alternative model is the following (I am abstracting away from the differences between the accounts of Doggett and Egan 2007, Currie and Ravenscroft 2002, and Velleman 2000 here): the pretend action is motivated by two different mental states—an imaginary version of a belief and an imaginary version of a desire. To use Currie and Ravenscroft's terminology, it is motivated by a "belief-like imagination" and a "desire-like imagination." So my daughter has a belief-like imagination that she is making pizza, and a desire-like imagination to spread the toppings evenly. And these two imaginary episodes motivate her pretend action.

Note the difference between this model and the one Nichols and Stich proposed. According to the Nichols and Stich explanatory scheme, the pretend action is motivated by an actual (conditional) belief and an actual desire. According to the alternative model, it is not motivated by either an actual belief or an actual desire, but by two imaginative episodes: a "belief-like imagination" and a "desire-like imagination."

Instead of focusing on the subtleties of the debate between these two explanatory schemes, I would like to point out an important assumption they have in common. An important common denominator between the two accounts of pretense is that they both take the belief–desire model of motivation for granted: real (not pretend) actions are triggered and motivated by beliefs and desires, and pretend actions are triggered and motivated either by beliefs and desires as well (as Nichols and Stich would argue) or by mental state types that are the imaginary versions of belief and desires (as the proponents of the alternative model would have it).

In Chapter 4 I argued against the belief–desire model as the general account of describing what triggers actions. Here I will argue against explanations of pretend actions in terms of the belief–desire model—both Nichols and Stich's and the alternative explanation. As with my attack on the belief–desire model in Chapter 4, my argument here has limited scope. There are real (non-pretend) actions that are in fact motivated by beliefs and desires. But not all actions are. And, similarly, there are pretend actions that can be

explained with reference to the imaginary versions of beliefs and desires. But not all pretend actions can be so explained.

Take the following pretend action. I pretend to raise a glass and take a sip from it, even though my hands are empty. What is the immediate mental antecedent of this action? That is, what representational state allows me to hold my hand and move it towards my mouth in the way I do? The two accounts of pretend actions I outlined above do not have a clear answer to this question. The belief (or the belief-like imagination) that I am making a toast does not specify what grip size I should maintain while I am performing the action. Nor does the desire (or desire-like imagination).

But we cannot rely on pragmatic representations either (whose job it would normally be to help us to have the appropriate grip size): they attribute action-properties to perceived objects, but there is no perceived glass—I am raising my empty hand. It seems that we need another kind of imaginative episode in order to explain what makes me hold my fingers the way I do, and what makes me move my hand the way I do. But this imaginative episode does not have propositional content like belief-like imagination or the possible world box. My proposal is that it is pragmatic mental imagery.

We have seen that both the Nichols and Stich account and the alternative account use a version of the belief–desire model, and modify it in such a way that it can explain pretend action. I am not using the belief–desire model. According to my account, the immediate mental antecedent of action is pragmatic representation. In the case of some (but not all) actions, this pragmatic representation is preceded and brought about by beliefs and desires. But pragmatic representation, unlike belief and desire, is necessary for the performance of an action.

If we apply this model to the explanation of pretend actions, what we get is that any of these three kinds of mental states can take an imaginary form. We have already seen that we can have belief-like imagination instead of beliefs, and desire-like imagination instead of desires. But, crucially, we can also have "pragmatic representation"-like imagination—or, as I call it, pragmatic mental imagery. Depending on which mental state takes on an imaginary version, we get very different kinds of pretense actions.

When I am pretending to raise my glass with nothing in my hands, I supposedly have belief-like imagination that I am drinking a glass of

wine (and maybe corresponding desire-like imagination), but in order for this belief-like imagination to have any influence on my actual movements, I also need to have pragmatic mental imagery that allows me to hold my fingers and move my hand in a certain way. This pragmatic mental imagery attributes action-properties (weight, shape, and spatial location properties) to the imagined glass in my hand. We have pragmatic mental imagery of the weight-property, the shape-property, the size-property, etc. of the nonexistent glass, and the attribution of these properties guides my pretend action: it guides the way I hold my fingers (as if around a glass), the way I raise my hand (as if raising a glass), etc. This pretend action cannot be explained without appealing to pragmatic mental imagery.

It is important to emphasize that I do not need to consciously visualize the glass in order to attribute various properties to it by means of mental imagery. Nor does this mental imagery need to be triggered voluntarily. But I do need to attribute action-properties to the glass by means of mental imagery—otherwise I would not know how to move my hand.

But isn't this too quick? Couldn't one argue that we do not need to appeal to pragmatic mental imagery here in order to explain our movements—we can just postulate a belief that would represent the ways in which I need to move my hand in order to perform this pretend action. In this case, the belief–desire model would need no readjustments in the case of pretense actions. The problem with this suggestion is that it leads to a significant structural asymmetry between real actions and pretense actions. We have seen at length that in the case of real actions, action-properties are not attributed by beliefs. So it would be natural to assume—given that the control of our bodily movements is presumably similar in the real and the pretense case—that in the case of pretense actions, action-properties are not attributed by beliefs either. The postulation that action-properties are attributed by beliefs in the case of pretense actions, but not in the case of real actions, is also difficult to reconcile with various empirical findings concerning the similarities (and differences) between the fine-tuned movements of the execution of the same action in real and pretense scenarios (see, for example, Cavina-Pratesi et al. 2011); as well as with cases of "semi-pretense" (Van Leeuwen 2011, see also Schellenberg forthcoming) where we have a mix of "real" and pretended actions.

In the example of pretending to take a sip from a nonexistent glass, my belief, my desire, as well as my pragmatic representation, are all "imaginary": belief-like imagination, desire-like imagination, and pragmatic mental imagery. But in other cases of pretend actions, our pragmatic representation is not imaginary: we do not need pragmatic mental imagery. Here is an example: I am taking a sip from a glass of cheap and bad red wine, and I pretend that I am taking a sip from a glass of 2004 Brunello di Montalcino. I may be using belief-like imagination (and, presumably, desire-like imagination), but the pragmatic representation that guides this action is exactly same as it would be if I were not pretending.[4] Pretense can happen without pragmatic mental imagery.

Can pretense happen without belief-like imagination, so long as what guides our action is pragmatic mental imagery? I am not sure. If we side with those (like Davidson 1980) who are reluctant to attribute beliefs (let alone belief-like imaginations) to animals, then, as it seems uncontroversial that at least some animals are capable of pretense (Mitchell 2002, Patterson and Cohn 1994, especially p. 285, see also Gómez 2008's cautious summary), pretense can happen without belief-like imagination. But even if we are willing to describe animals as having beliefs, maybe attributing belief-like imagination to them would amount to anthropocentrism or over-intellectualization. But as animals must have pragmatic representations (otherwise they could not perform goal-directed actions: chase prey, escape from predators) and, as we have seen, they also have mental imagery (even pigeons, see Rilling and Neiworth 1987, Neiworth 1992, see also Oakley 1985 for a summary), we have good reason to assume that they can have pragmatic mental imagery,[5] which could be used to explain their pretense behavior. The same goes for infants.

---

[4] Note that in this case the actual *basic* action and the pretend *basic* action are the same: taking a sip from a glass. But the non-basic actions are different. Hence, as we have seen, although the pragmatic representation that makes the pretend action possible is the same as the pragmatic representation that makes the actual action possible, the attributed thick action-properties can be, and presumably are, very different. So one way of describing the example is that as a result of the belief-like imagination, we attribute imaginary thick action-properties.

[5] The recent findings that animals, even birds, are capable of episodic memory—memory a distinctive feature of which is that it is accompanied by mental imagery—suggests the same. See Clayton et al. 2001 and Emery and Clayton 2004.

This leads to a non-monolithic account of pretense actions. Some of our pretense actions can be explained with the help of belief-like (and desire-like) imagination. Some can be explained with the help of pragmatic mental imagery. And in some cases, we need to appeal to both kinds of imaginary states.

Neil Van Leeuwen recently criticized both Nichols and Stich's theory and the alternative account (Van Leeuwen 2011). One of his main objections is that neither theory can account for an interesting special case of pretense that he calls "semi-pretense"—when we are both pretending and performing a real action. His example is a scenario where two kids are watching some other kids who are jumping off the high-dive, and evaluating the quality of their dive by holding up their fingers. Is this a real action? In some sense, it is: they are genuinely evaluating the dive of their friends. But it is also pretense inasmuch as they pretend to be judges who give points for each dive. It is not full pretense though—they do not hold up placards with numbers, they merely use their fingers. Van Leeuwen argues that the propositional accounts cannot explain semi-pretense because there is no middle ground between the possible world box and the belief box (there is no semi-possible world box), and there is no middle ground between belief and belief-like imagination.

My account can explain "semi-pretense" because it allows for the "integration of perception and imagination," the importance of which Van Leeuwen is right to emphasize (Nanay 2010c and Nanay 2007 do the same in a different context—in the context of amodal perception). Sometimes, for example when I am pretending to stab you with a sword with my hands empty, I attribute all the relevant properties to the non-existent sword by means of pragmatic mental imagery. Some other times—for example, when I am pretending to stab you with a sword and I in fact hold an umbrella in my hand—I attribute some of these properties perceptually (its weight, for example), but I attribute others by means of mental imagery (for example, the property of where the end of the sword is and how sharp it is). In this case, some of the action-properties are attributed by my pragmatic representation, and some others by my pragmatic mental imagery. Van Leeuwen's main reason for discarding Nichols and Stich's account and the alternative account was that they could not explain this integration between perception and imagery (and as a result, they could not explain "semi-pretense"). My account can.

Van Leeuwen's positive account is, in some respects, similar to my own. Importantly, he also argues that the imaginative episode that explains pretense action is not a propositional attitude, but more akin to mental imagery. Further, he also emphasizes the importance of action in this imaginary episode. It may be worth emphasizing the differences between his account and mine—noting that these differences are minor in comparison with the difference between the propositional accounts on the one hand and Van Leeuwen's and my account on the other.

According to Van Leeuwen, there are two imaginative episodes involved in bringing about a pretense action. The first one is genuine mental imagery. This has no action-guiding or action-oriented characteristics, but it gives rise to another imaginative episode that is action-guiding and action-oriented. It is this second imaginative episode that serves as the immediate mental antecedent of our pretense actions. The relation between these two imaginative episodes is complex: the second (action-oriented) one can, but does not have to, take the first (mental imagery) one as its content, for example. The main difference between this account and my own is that I only appeal to one imaginative episode: pragmatic mental imagery. Pragmatic mental imagery is both genuine mental imagery (as Van Leeuwen's first imaginative episode), and the immediate mental antecedent of many of our pretense actions (like Van Leeuwen's second imaginative episode). Further, we have independent reasons to postulate this kind of mental state.

5.5 Aliefs

Tamar Szabó Gendler argues in a series of influential articles for the importance of a mental state type she calls "aliefs" (Gendler 2008a, 2008b, 2011, see also McKay and Dennett 2009, Bloom 2010). An alief is "an innate or habitual propensity to respond to an apparent stimulus in a particular way" (Gendler 2008a, p. 553). Aliefs, unlike beliefs, are not under our conscious control. We can, and very often do, have a belief that x is F and an alief that x is not F at the same time.

Since Gendler introduces the concept of alief by example, I will consider all of the examples she provides:

(a) "A frog laps up the BB that bounces past its tongue" (Gendler 2008a, p. 552).

(b) "A puppy bats at the 'young dog' in the mirror in front of him" (Gendler 2008a, p. 552).

(c) "A sports fan watching a televised rerun of a baseball game loudly encourages her favourite player to remain on second base" (Gendler 2008a, p. 552).

(d) "A cinema-goer watching a horror film 'emits a shriek and clutches desperately at his chair'" (Gendler 2008a, p. 553, see also Gendler 2008b, p. 637).

(e) "A man suspended safely in an iron cage above a cliff trembles when he surveys the precipice below him" (Gendler 2008a, p. 553, see also Gendler 2008b, pp. 634–5).

(f) "An avowed anti-racist exhibits differential startle response when Caucasian and African faces are flashed before her eyes" (Gendler 2008a, p. 553).

(g) "A person who has set her watch five minutes fast [rushes], even when she is explicitly aware of the fact that the time is not what the watch indicates it to be" (Gendler 2008a, p. 555).

(h) "[A person feels] reluctance to eat fudge shaped to look like dog feces" (Gendler 2008a, pp. 555–6, see also Gendler 2008b, pp. 635–6).

(i) "[A person feels reluctance] to drink lemonade served in a sterilized bedpan" (Gendler 2008a, p. 556, see also Gendler 2008b, pp. 635–6).

(j) "[A person feels reluctance] to throw darts at a picture of a loved one—even when she explicitly acknowledges that the behaviors are harmless" (Gendler 2008a, p. 556, see also Gendler 2008b, pp. 635–6).

(k) "[A person feels] hesitant to sign a 'pact' giving her soul away to the devil—even if she is an atheist, and even if the pact says explicitly at the bottom 'this is not a real pact with the devil; it is just a prop in a psychology experiment'" (Gendler 2008a, p. 556).

(l) "The Hitchcock expert [has a tendency] to experience suspense as the shower scene proceeds, even though she has written a book detailing *Psycho* frame-by-frame" (Gendler 2008a, p. 556).

(m) "A chef who has recently rearranged his kitchen [has the tendency] to walk towards the old knife drawer to get his cleaver, even as he talks about how happy he is with the new set-up" (Gendler 2008a, p. 556).

(n) "Subjects whose aim is to select a red ball [have the propensity] to go with frequency (choosing from a bag with 9 red and 91 white balls) rather than probability (choosing from a bag with 1 red and 9 white balls)—even when the comparative likelihoods are prominently displayed" (Gendler 2008a, p. 556).

These are diverse cases, and Gendler is quick to emphasize that aliefs are not supposed to be part of some kind of final theory of how the mind works. Her claim is much more modest: if we talk about beliefs and desires, we also need to talk about aliefs. This is an important contrast between Gendler's account and the "associative–propositional evaluation" model (Gawronski and Bodenhausen 2006, Gawronski et al. 2009) with which it shares some important features (see Nagel forthcoming for a comparison).

Some of the characteristics Gendler attributes to aliefs are strikingly similar to the characteristics of pragmatic representations. She says that aliefs are "Associative, automatic, arational, shared by human and non-human animals, conceptually antecedent to other cognitive attitudes that the creature may go on to develop, action-generating and affect-laden" (Gendler 2008a, pp. 557–8).

The aim of this section is not to look for minor disagreements between Gendler's account and mine (there are many of these),[6] but to argue that the conceptual framework of pragmatic representations and pragmatic mental imagery can elucidate many (maybe most) of the examples of aliefs (but not all of them). Gendler says that she is "fully open to the possibility that [she has] misdrawn the boundaries of the mental state that [she is] interested in" (Gendler 2008a, p. 555). I would like to propose exactly this kind of (slight) redrawing of boundaries in what follows.

---

[6] Two quick examples. First, Gendler seems to presuppose that while beliefs can and do change frequently (when updated in the light of further evidence), aliefs are more stable and less likely to change. As we have seen in Chapter 2 with the basketball example, often the non-conscious mental state (the pragmatic representation) can change, while every belief/thought/experience remains the same (there are many other examples; see, for example, Olson and Fazio 2006). Second, Gendler claims that beliefs are more likely to track truth (because they respond to rational revision). But, as the three-dimensional Ebbinghaus experiment I considered in Chapter 2 shows, this is not always the case: our conscious experiences and beliefs often mislead us, while our unconscious pragmatic representations can track the real properties of objects (again, there are other examples; see, for example, Bechara et al. 1997, 2005).

A tempting reaction in the face of the postulation of aliefs is skepticism: why should we postulate a new kind of mental state? After all, the only feature the examples of aliefs have in common is that in all of them there is some kind of unconscious influence on our behavior. What this unconscious influence may be can vary widely: we do not need to postulate a new mental state type; there may be many different kinds of unconscious mental state types that could influence our behavior. I want to resist this tempting reaction by appealing to pragmatic representation and pragmatic mental imagery that, taken together, cover almost all the examples aliefs are supposed to cover. In short, Gendler put her finger on a very important phenomenon, but I want to argue that we can give a more thorough and detailed explanation for this phenomenon with the help of pragmatic representations and pragmatic mental imagery.

Let us go back to Gendler's examples for aliefs. Some of them sound very much like examples of pragmatic representations. Examples (a), (b), and (e) are clearly among them. When the frog snaps at the BB, it has to represent it as being at a certain spatial location, otherwise it could not stick out its tongue in the right direction. In other words, it must have a pragmatic representation about the whereabouts of this object that would guide its snapping movement. The puppy, similarly, must have a pragmatic representation that allows it to bark at the "other puppy." Assuming that the dog is conscious, it also attributes thick action-properties to the "other puppy": maybe the property of being scary or of being a potential playmate, etc. And the man in the cage over the precipice has a pragmatic representation that is clearly insensitive to his beliefs about the safety standards of his cage.

In all these cases, the phenomenal character of alief is colored by the attribution of thick action-properties that accompany our pragmatic representations. As we have seen, pragmatic representations attribute action-properties, but when we have pragmatic representations we can also be aware of a wider range of thick action-properties, properties like edible, being an obstacle, and so on. Thick action-properties are relational properties that cannot be fully characterized without reference to the agent's action, and they are often emotionally charged. The dread the man in the cage feels can be explained with the help of the thick action-properties he is aware of.

Pragmatic representations are perceptual states and they are, in some ways (not in others, see the end of Chapter 3), informationally encapsulated from the rest of the mind—as we have seen in relation to the three-dimensional Ebbinghaus illusion and the basketball/goggles example in Section 2.5. But then pragmatic representations satisfy the main requirement for aliefs: namely, that they influence our behavior even though this goes against everything we explicitly believe. My behavior when picking up the chip in the three-dimensional Ebbinghaus illusion scenario is very much influenced—in fact, it is guided—by my pragmatic representation, but this goes against everything I believe about the scene in front of me. The same goes for the basketball/goggles example: my beliefs do not change at all. But my behavior does change as a result of a change in my pragmatic representations. Pragmatic representations can explain some of the examples Gendler describes as aliefs.

But many, even most, of the examples for aliefs are not pragmatic representations: in these cases, we do not see anything as having action-properties. My claim is that the vast majority of these examples can be explained if we appeal to pragmatic mental imagery.

In (m), the chef forms a pragmatic mental imagery of his cleaver as being in the old drawer because he has seen it there very often in the past. And that is why he goes to that drawer in spite of having an explicit belief that the cleaver is not there. His pragmatic mental imagery attributes very specific spatial location properties to the cleaver—and ones that are very different from the spatial location properties he believes it to have. And it is the attribution of these action-properties (i.e., the spatial location properties) that guides the chef's action. As we have seen, mental imagery in general, and pragmatic mental imagery in particular, can be involuntary (remember the earworms)—it can go against your beliefs and desires, and this is exactly what happens in this case.

Note the structural similarity between this example and the basketball/goggles scenario from Chapter 2. In both cases, there is a discrepancy between the properties attributed consciously (by our beliefs) and the properties attributed by our pragmatic representation/pragmatic mental imagery. And in both cases, it is the pragmatic representation/pragmatic mental imagery that influences the performance of our action.

A slightly different explanatory scheme works for (i) and (h): when we see the sterilized bedpan full of lemonade and we are told to drink it, we

can't help but form a pragmatic mental imagery of the bedpan being full of urine, and this is what makes us hesitate to drink it. Remember that we can form pragmatic mental imagery of an object that we do perceive—when we attribute imaginary action-properties to the perceived object. And this is what happens here. We see the bedpan full of lemonade, but can't help forming the pragmatic mental imagery of the bedpan full of urine (presumably as a result of simple associative perceptual learning). Similarly, when we see the fudge shaped to look like dog feces that we are supposed to eat, this leads to the pragmatic mental imagery of dog feces. Here, what makes the action of eating the fudge not particularly appealing is not as much the attribution of action-properties, but of the emotionally colored thick action-properties we experience the imagined dog feces as having. But, as we have seen, pragmatic mental imagery, if conscious, is normally accompanied by the awareness of a rich set of thick action-properties.

Example (f) can be analyzed in a similar manner. The avowed anti-racist forms different pragmatic mental imagery in response to seeing an image of a Caucasian and an African face—the action-properties, and especially the thick action-properties, attributed by this pragmatic mental imagery are different. And this explains the difference in her startling. This way of analyzing the example would also explain the following experimental finding: if the task is to detect the presence or absence of dots on a photo of a face, the automatic stereotype activation effect when seeing minority faces is eliminated (Macrae et al. 1997). If the automatic stereotype activation effect has to do with pragmatic mental imagery, then a motor task that involves pragmatic representation should be expected to diminish this effect.

The real question is, of course, why does she form this pragmatic mental imagery on the basis of seeing the two different faces? And here some recent empirical findings point towards perceptual learning (Olson and Fazio 2006), which supports my proposal that what is responsible for this difference in the startling reaction is something perceptual/quasi-perceptual: pragmatic mental imagery.

How about case (j), which is based on studies in Rozin et al. (1986)? Subjects are reluctant to throw darts at the picture of a loved one, and if they do throw darts they are more likely to miss than if they are throwing darts at a picture of someone they do not know. There are two ways of analyzing this, depending on one's views on picture perception. One

possibility is to argue that in order to explain this phenomenon we need to appeal to pragmatic mental imagery. The general proposal would then be that when we see the picture of the loved one, we can't help forming a mental imagery of the actual person (as opposed to the picture of her).[7] And this involuntary mental imagery of the loved person (as opposed to the picture itself) attributes action-properties, and presumably also thick action-properties, like "not suitable for throwing darts towards." In other words, on the basis of the perception of the picture of the loved one in this situation, we form an involuntary pragmatic mental imagery of her, and it is this pragmatic mental imagery, not the perception of the picture, that makes us reluctant to throw the darts.

But this is not the only way in which we can analyze this case in my framework. Those who do not like the appeal to mental imagery when talking about picture perception can go another way—they can describe this scenario with the help of straightforward pragmatic representations. If we do indeed perceive things in pictures (Wollheim 1980, 1987, 1998), then we can have a perceptual state of the depicted object—in this case, our loved one—that attributes action-properties to them: this is a pragmatic representation, not a pragmatic mental imagery. Then we can explain the belief-insensitivity of our movement with reference to the belief-insensitivity of pragmatic representations—like in case (b), where the puppy is barking at its own mirror image.[8] We attribute action-properties to the depicted object we see in the picture—not to any imagined object.

The phenomenon Gendler describes with example (d) is sometimes analyzed as the "paradox of fiction": how is it possible to have a strong emotional reaction to something we believe to be fictional? Without trying to give a general answer to the paradox of fiction (something

---

[7] It is important to note that with this claim one does not necessarily side with Waltonian accounts of depiction, according to which when we see the picture of X we imagine our experience of the picture to be of X (Walton 1990, 2002, Maynard 2012). But if one is tempted by the Waltonian account, this may be a natural way of extending that account to these cases. My own account of depiction is very different from Walton's (see Nanay 2004, 2005, 2008, 2010f, 2011c).

[8] Again, in terms of theories of depiction, this way of analyzing the dart-throwing example does not necessarily lead to an account of picture perception as "seeing-through." While some people have argued that the perception of at least some pictures is like seeing through a mirror (Lopes 2003), we do not have to endorse this in order to analyze the dart-throwing scenario in terms of pragmatic representations.

I attempt to do elsewhere, see Nanay forthcoming a), the specific example Gendler alludes to can be explained if we appeal to pragmatic mental imagery or to pragmatic representations—depending on one's preferred account of picture perception. What makes the cinema-goer "clutch desperately at his chair" is the attribution of action-properties as well as emotionally very salient thick action-properties to the monster (to the depicted monster we literally see in the picture or to the imagined monster we form our pragmatic mental imagery about). And it is this attribution of thick action-properties that explains our reaction.

It would be tempting to use the same explanatory scheme for example (l): the Hitchcock expert watching the shower scene in *Psycho*. But I think that this case is even more complicated: we do attribute thick action-properties when we see this scene, but this attribution of thick action-properties constitutes a distinctive and important case—we see the murderer as relevant not to our own actions, but to the actions of the woman in the shower. One way of putting this would be to say that we see the murderer as "affording an action"[9] to the woman in the shower. Not to ourselves, but to someone else. This is still an instance of the attribution of thick action-properties, but these thick action-properties are not "self-centered" but "other-centered." I call this phenomenon "vicarious perception," and spend Chapter 6 explicating it.

For now, it is enough to note that the attribution of other-centered thick action-properties is (almost) as automatic and insensitive to our beliefs and thoughts as the attribution of self-centered thick action-properties. Thus, vicarious perception is another way in which we can flesh out some examples of aliefs. The example of the sport fan (example (c)) can also be analyzed in terms of vicarious perception. We see the "third base" as "affording an action" to the baseball player, and this mental episode is insensitive to our explicit belief that we are watching a recording of the baseball game that was played hours ago.

I said that many, maybe most, examples of alief can be analyzed as pragmatic representations or pragmatic mental imagery (or vicarious perception). My claim is not that all examples can be. I am not sure about (g), (k), and (n). In these cases, whatever is responsible for the behavior seems to be a relatively complex mental state. It may not be impossible to

---

[9] More on what I mean by this in Chapter 6.

stretch my explanatory scheme to cover these cases. But the aim of this chapter was not to explain all cases of aliefs in terms of pragmatic representation and pragmatic mental imagery, but to show that many examples that motivated Gendler to posit aliefs can be explained if we appeal to mental states that we have independent reasons to posit: pragmatic representations and pragmatic mental imagery.

Gendler briefly discusses the possibility of explaining alief in terms of imagination (Gendler 2008a, pp. 567–8), but she concludes that while, at least in some cases, "imagination gives rise to behavior via alief" (Gendler 2008a, p. 568), aliefs themselves cannot and should not be identified with imagination. My attempt was to use a very specific kind of imagination, mental imagery, and even a very specific kind of mental imagery, pragmatic mental imagery, in order to explain at least some of the examples Gendler explains in terms of aliefs.

Gendler explicitly states that alief is not a "fundamental mental category, one that will be part of our 'final theory' of how the mind makes sense of the world" (Gendler 2008a, p. 555). In some ways, I take myself to be continuing Gendler's project. She identified a crucial phenomenon, and I am trying to explain it in terms of pragmatic representations and pragmatic mental imagery: mental states that we have independent reasons to posit.

6

Vicarious Perception

Virginie's face became her own, Virginie's dress clothed her, Virginie's heart was beating in her breast.

Gustave Flaubert, *Un Coeur Simple*, 1876.

6.1 Vicarious perception

The first four chapters of this book were about pragmatic representations: the perceptual attribution of action-properties to an object. In Chapter 5, I considered an interesting and important variant of pragmatic representation: the attribution of action-properties to an imagined object (or of imaginary action-properties to an object)—pragmatic mental imagery. Both of these mental processes represent (perceived or imagined) objects as relevant to the performance of one's own action. This chapter considers another interesting and probably even more important phenomenon in the vicinity: when you see something as relevant to the performance of not your own action, but to the action of someone else. I call this phenomenon vicarious perception.

Pragmatic representations attribute action-properties that could be called "self-centered": properties that need to be attributed in order *for me* to perform an action. As we have seen, when we have a pragmatic representation of an object, we also tend to experience it as having "thick action-properties": properties that can be richer than action-properties proper. But these thick action-properties are also "self-centered": they pertain to my performance of the action.

Vicarious perception attributes "other-centered" action-properties and "other-centered" thick action-properties that pertain to someone else's performance of an action. Self-centered action-properties are the properties the representation of which is necessary for the performance

of *my* action, whereas other-centered action-properties are the properties the representation of which is necessary for the performance of *someone else's* action. Self-centered thick action-properties are relational properties that cannot be fully characterized without reference to *my* action, whereas other-centered thick action-properties cannot be fully characterized without reference to *someone else's* action.

We have seen that "being edible" and "being climbable" are thick action-properties. When I see a very tasty-looking cake, I may attribute the thick action-property of being edible (for me). When I want to climb a tree, I may attribute the thick action-property of being climbable (for me). These are self-centered thick action-properties. In vicarious perception, we attribute other-centered thick action-properties: I see the tree as climbable not for me but for you, and I see the cake as edible not for me but for you.

To use a metaphor I have been avoiding so far, I see the tree and the cake as affording a certain action (of climbing or eating) not to myself, but to another agent. I have been avoiding this term of "affording an action" in the previous chapters because it may create some confusion—and may tempt some to interpret my views in a Gibsonian light, which, as I argued at length in Chapters 1 and 3, is in some ways the opposite of what I am trying to do. But I will use the concept of "affording an action" in what follows, merely because it is convenient to use. However, when I say that someone represents an object as affording an action, all I mean is that someone represents this object as having a thick action-property. And when I say that someone represents an object as affording an action to someone else, all I mean is that someone attributes an other-centered thick action-property to this object. When we perceive vicariously, we see things as affording actions to someone else. We attribute other-centered action-properties to them.

As we have seen in Chapter 3, while action-properties are perceptually attributed, I want to remain neutral about whether thick action-properties are necessarily perceptually attributed (because of the difficulties of drawing a clear distinction between perceptual and non-perceptual phenomenology). As we attribute self-centered action-properties perceptually, we should expect that we also attribute other-centered action-properties perceptually. However, as the attribution of self-centered thick action-properties is not necessarily perceptual, the attribution of other-centered thick action-properties is not necessarily perceptual either.

When I talk about "seeing objects as affording actions to someone else" and about "vicarious perception," this then does not strictly speaking entail that the property of affording actions to someone else is represented by a perceptual state. I nonetheless use the terminology of "vicarious perception" in order to indicate the simplicity of these processes. Purists may want to read "vicarious quasi-perception" in its stead.

Here is an example. I am sitting in my armchair looking out of the window. I see my neighbor running to catch her bus. There are lots of people in the street and my neighbor is zigzagging around them on her way to the bus that is about to leave. How will I represent the lamppost in my neighbor's way? I will not see it as affording an action to me: I am not about to perform any action, let alone an action that would involve the lamppost. But I don't see it in a detached, action-neutral way either: I see it as affording an action (of bumping into) to my neighbor. I represent the lamppost as having a thick action-property, but this thick action-property is a relational property that is determined by the object and my neighbor, rather than by the object and myself.

Another, somewhat more evocative example: I am still sitting in my armchair looking out of the window. I see my neighbor standing in the middle of the street, deep in thought, and I also see that unbeknownst to her a tiger is approaching from behind and is about to attack her. Again, I represent the tiger in an action-oriented manner—not as affording an action (of running away, presumably) to myself, but to my neighbor. This example will play an important role in the next section.

In this example, I have no idea about what specific action the tiger affords to my neighbor. Is it running away? Is it turning around to face the tiger? Is it shooting it with a tranquilizer? In some other cases, however, one may have a very clear idea about the specific action something affords to another agent. Suppose that you are looking at your child carrying a full glass of water very carefully from the kitchen to the living room. You see that the glass is slightly tilted—the water is about to be spilt. You attribute a very specific other-centered action-property to the glass: you see it as affording a very specific action to your child. In this case, the property you attribute is an other-centered action-property: a property the representation of which is necessary for the performance of your child's action.

In the tiger example, in contrast, I attribute an other-centered *thick* action-property: the property I attribute to the tiger is not a property the

representation of which is necessary for the performance of my neighbor's action—it is not clear how my neighbor's action should be specified to begin with. But it is a property that is not possible to fully characterize without reference to *some* action of my neighbor's.

Remember the example of trying to diffuse a bomb. You have no idea how to do it; hence, you are not in a position to attribute action-properties to it. But you still experience the bomb as having *thick* action-properties. Thick action-properties, as we have seen, specify actions on a much more general level than action-properties proper. The representation of action-properties proper needs to be able to guide our very specific actions. There is no such constraint on the representation of thick action-properties.

The tiger example is the equivalent of the bomb-case in an other-centered context. You are not in a position to attribute action-properties proper, but you do experience the tiger as having thick action-properties: properties that cannot be fully characterized without reference to some action of your neighbor's. Both the attribution of other-centered action-properties proper and the attribution of other-centered thick action-properties count as vicarious perception.

We perceive things vicariously all the time. For example, when we are following a sport event on TV—say, a football game—we see the ball as affording various actions for a certain player. We also often perceive vicariously when we are watching movies. Suppose that you are watching a Tom and Jerry cartoon, where Tom is chasing Jerry, who gets stuck in a corner where a hammer happens to be lying around. You see the hammer as affording a certain action not to yourself, but to Jerry (or to Tom, depending on who you're identifying with). More generally, the concept of vicarious perception can also help us understand what happens when we identify with a fictional character when we are watching a film or attend a theater performance (Nanay 2006b, forthcoming a).

I will argue that vicarious perception is a very basic form of social cognition: a very basic way of engaging with others. It can play an important role in a number of debates in cognitive science: about non-human social cognition, about the cognitive development of human children, and about joint action. And it can also help us explain some rudimentary forms of social emotions.

6.2 Vicarious perception versus theory of mind

I will argue that the theoretical framework of vicarious perception can be used as a promising alternative to the existing theoretical framework of "theory of mind" in philosophy of mind and the cognitive sciences. First I need to clarify how the concept of "theory of mind" has been used, and what theoretical role it is supposed to play in philosophy of mind and the cognitive sciences.

6.2.1 Theory of mind

We attribute mental states to others all the time: if I know that my wife wants diamonds for her birthday, this may influence my actions. And if I know that my daughter knows how to open her bedroom door, this may also influence my actions. We use the ability of attributing mental states—beliefs, desires, wishes, etc.—to others with great ease and frequency. In fact, it would be very difficult to get by without this ability, which is known as "theory of mind."

In the last three decades, there has been a tremendous amount of research in contemporary philosophy of mind, psychology, and primatology on "theory of mind." More specifically, this research has been mainly focused on the following four questions:

(i) Do non-human animals have theory of mind?
(ii) How does theory of mind develop in ontogeny?
(iii) What mental processes make theory of mind possible in adult humans?
(iv) What are the neural underpinnings of theory of mind?

How we answer these questions clearly depends on how we define "theory of mind." Above, I used the term "theory of mind" to refer to the ability to attribute mental states to others, and this is indeed the original characterization given in the paper that introduced this concept (Premack and Woodruff 1978). But it needs to be noted that there is no consensus about the exact meaning of the term (see Whiten 1994). Some talk about "theory of mind" as the "capacity to reason about mental states" (Penn and Povinelli 2007, p. 731), some others as the ability to understand others' mental states (Call and Tomasello 2008). Peter Carruthers and Peter K. Smith, in their introduction to the most

important edited volume on "theory of mind," take it to mean possessing "a conception of the mind of another creature" (Carruthers and Smith 1996, p. 1). Cecilia Heyes, in an equally influential survey article, claims that "an animal with theory of mind believes that mental states play a causal role in generalizing behavior and infers the presence of mental states in others by observing their appearance and behavior" (Heyes 1998, p. 102, see also Heyes 1994). Others hold a somewhat similar view, according to which only passing the false belief test provides evidence for theory of mind (Gómez 2004).

The differences between these formulations are significant: Heyes's concept is quite demanding—it requires the animal to represent a causal relation between an unobservable and an observable entity. In comparison, the original Premack and Woodruff concept is more liberal: attributing a mental state may not even require the representation of another representation (see Perner 1991).

It is not entirely clear what the key concepts in any of these characterizations are supposed to mean: what is it to "understand" someone else's mental states? Does understanding imply knowledge? If we think of "theory of mind" as a capacity to reason about others' mental states, is this supposed to be theoretical or practical reasoning? Further, it is not clear what is meant by the mental states that are being attributed or understood—would any mental state do? If a creature can attribute perceptual states to others but not beliefs, are we justified in describing her as having "theory of mind" (Bermúdez 2009)?

Recently, there has been a move away from the use of the concept of theory of mind, especially in the developmental psychology literature. But the proposed alternatives do not seem to have much advantage over the original concept. For example, the concept of "perspective taking" (Moll and Tomasello 2006, Tversky and Hard 2009) is not itself well defined: it is not clear, for example, whether the perspective in question is epistemic or spatial (or both).

These terminological confusions have made it difficult to settle questions (i)–(iv). Take (ii), for example. Depending on which concept of theory of mind we are using, we get different answers about its ontogeny, and, importantly, about the age at which this capacity appears. If we equate theory of mind with the ability to attribute *false* beliefs (Wimmer and Perner 1983), then theory of mind appears at a later age than it does if we interpret it as "taking the intentional stance" (Gergely et al. 1995).

Or, to take (i), if we accept Heyes's concept of theory of mind it is likely to rule out all non-human primates from the elite circle of creatures with theory of mind. But if we think of theory of mind as the attribution of perceptual states (Hare et al. 2000, 2001), we may be justified in talking about primate theory of mind (for various potential problems with this approach, see Section 6.3).

It would be easy to dismiss these divergences in the use of the concept of theory mind as irrelevant as there is one *correct* concept of theory of mind that latches on to a real mental capacity, whereas all the other interpretations are just wrong. I think the situation is more complicated. The emphasis on one single concept that explains much of social cognition—namely, that of "theory of mind"—may have been a methodological mistake.

Social cognition is not a monolithic category, and "theory of mind" is only one, relatively complex, form of making sense of others. It does not help us if we want to understand the origins (phylogenetic or ontogenetic), and some more rudimentary forms, of social cognition. I argue that empirical findings from developmental psychology and primatology point to a possible alternative: vicarious perception. If we take questions (i)–(iii) (I will not say much about (iv), although it is a thriving research program; see, for example, Siegal and Varley 2002, Gobbini et al. 2007, for summaries) to be about vicarious perception, and not about "theory of mind," we get non-trivial answers that are clear and unambiguous.

6.2.2 Vicarious perception is not theory of mind

Vicarious perception does not count as theory of mind in any of the definitions of theory of mind that I mentioned above. Importantly, I can engage with someone else this way without attributing any mental (or even perceptual) state to her. In the tiger example, I know that my neighbor does not see the tiger, but I nonetheless see the tiger as a threat to her.

In fact, our cognitive (and maybe emotional) engagement with others is as strong (or maybe even stronger) if there is an epistemic asymmetry between the two agents. If I know more than you do—if I see something close to you as a threat, whereas you don't—I am likely to react as strongly as (or maybe even more strongly than) I would if you were also aware of the threat. This very basic visceral "Oh my god!" or "Watch out!" reaction is difficult to explain in terms of the attribution of mental

states. Defenders of the idea of theory of mind could say that I attribute the mental (or maybe perceptual) state to you that you are *not* aware of the threat, and that I compare this mental state (attributed to you) with my own mental state of being aware of the threat, and that this comparison triggers my reaction. But this explanatory scheme presupposes extremely complex mental processes in order to explain a very simple and instinctive reaction.

My explanation is simpler: our visceral "Oh my god!" or "Watch out!" reaction is triggered by vicarious perception: we see the object as having other-centered thick action-properties, and highly emotionally charged ones—no attribution of mental states is needed to explain this.

Theory of mind entails some kind of representation or understanding of another person's mental state. Vicarious perception only presupposes a perceptual state—no complex representation of representation is required. In short, vicarious perception is very different from, and much simpler than, theory of mind.

To sum up, vicarious perception is the perceptual representation of an object as having some properties, and some of these properties are that it affords an action to another agent. This is not the attribution of a mental state to another agent, but the attribution of a property to an object— very different kinds of attribution indeed.

But the deeper question is whether seeing an object as affording an action to another agent *presupposes* the attribution of a mental state to another agent. My answer is to concede that sometimes it does: sometimes our vicarious perception presupposes theory of mind. But, crucially, this is not always the case. There are instances of vicarious perception (the tiger example above is one) that do not presuppose theory of mind. To be sure, vicarious perception does presuppose the attribution of agency. There is no vicarious perception if I see the tiger attacking a trashcan: I don't see the tiger as affording an action to the trashcan. But it doesn't presuppose the attribution of *mental states* to an agent. Vicarious perception is a way of engaging with someone else. This capacity, like theory of mind, is made possible by our much simpler capacity to identify some entities, but not others, as potential agents (see, for example, Scholl and Tremoulet 2000). But this does not make vicarious perception itself entail the representation of someone else's mental state.

We can, of course, have both. Theory of mind—the attribution of a mental state to someone else—can and often does influence our vicarious

perception. To use the tiger example once again, if I knew that my neighbor worked as a lion-tamer in a circus and was not at all afraid of tigers, my vicarious perception of the situation would be very different. Just as pragmatic representations are in certain contexts sensitive to top-down influences (remember the matchbox experiment from Chapter 2), vicarious perception can also be influenced by top-down effects. In fact, this happens very often. Suppose that I am watching my friend look for her glasses and that I can see where the glasses are, but she can't. If I attribute to my friend the desire to find her glasses, this mental state attribution will likely influence my vicarious perception of the glasses as affording an action to her.

It is important to emphasize that vicarious perception *can* entail the attribution of a mental state to someone else. The example of watching your child carry a full glass of water with great concentration may be a good example. In this case, seeing the glass as having action-properties may entail attributing a pragmatic representation to your child. I want to remain non-committal about this. But it is important to emphasize that other instances of vicarious perception do *not* entail the attribution of a mental state to the other agent. If your child does not pay attention to how she holds the glass and the water is about to get spilt, then if you see the glass as affording an action to your child, this does *not* entail the attribution of a pragmatic representation to her (she does not have any pragmatic representation involving the glass, after all). The same goes for the tiger example.

Recent empirical findings support my emphasis on the action-oriented nature of some basic forms of our cognitive engagements with others. It has been pointed out that we (adult humans) are much more likely to spontaneously describe things from someone else's perspective when the performance of an action is involved in some way (Tversky and Hard 2009). If we stick to the theory of mind framework, this result seems puzzling: if our engagement with others is all about attributing mental states to others, why would scenarios involving actions (in various ways) make this shift of perspective more likely? If we accept, however, that one important form of social cognition is vicarious perception, then these results are exactly what we should expect: we shift perspective when we see objects around us as affording actions to others.

Even more importantly, there are bimodal neurons in the ventral intraparietal area (VIP) of the macaque brain that are sensitive both to stimuli in the peripersonal space of the subject itself, and stimuli in the peripersonal space of other agents the subject sees (Ishida et al. 2009, see also Costantini et al. 2010, 2011, Cardellicchio et al. 2013). To put it very simply, these bimodal neurons are sensitive to both objects within one's own reach and objects within another agent's reach, making them a good candidate for one possible neural mechanism that may be responsible for vicarious perception. (Note that these bimodal neurons are very different from the ones the "mirror neuron" literature talks about, as we shall see in Section 6.2.3.)

6.2.3 Vicarious perception versus some recent alternatives to theory of mind

I have contrasted vicarious perception with theory of mind. But I also need to contrast it with two other proposals about social cognition.[1] The first suggestion is that the most basic form of social cognition involves merely seeing the mental states (or emotions) of another person on her face (Gallagher 2001, 2005, 2008, Zahavi 2008, Ratcliffe 2007, Hutto 2007, 2011, de Bruin, Strijbos and Slors 2011). There is no need to attribute mental states to others: the mental states of others are directly perceivable.

This proposal is similar to mine inasmuch as it also emphasizes the importance of perceptual processes in the basic forms of social cognition, but this similarity will turn out to be quite superficial. I am not sure how this account can be made more precise: what kind of perceptual process would count as the direct perception of someone's emotions or mental states? But whatever this perceptual process may be, it seems to presuppose that I see the other person's face (or at least her expressive bodily comportment). Note that there is no such requirement in the case of vicarious perception. If you are facing away from me, and I see a tiger attacking you from behind, we still get an instance of vicarious

---

[1] I do not mean to suggest that all recent accounts of social cognition belong to one of these three views (theory of mind, direct perception, mirror neurons). First, some have tried to combine the advantages of (a version of) the theory of mind approach and the mirror neuron approach (see, for example, Goldman 2006). Second, there are views that do not naturally fall under any of these three labels, for example the "minimal theory of mind" approach of Apperly and Butterfill 2009 and Butterfill and Apperly forthcoming.

perception. In the case of vicarious perception, there is no need for me to see your face (or even your "expressive bodily comportment"). All I need to see is the object that I represent as affording an action to you.

The direct perception of other people's mental states (if we can make this concept more precise) may or may not be one of our ways of engaging with others. But it is very different from vicarious perception.

The third proposal I need to contrast vicarious perception with (besides theory of mind and direct perception) is the recently popular account of social cognition in terms of mirror neurons (Gallese 2007, Gallese and Goldman 1998, Gallese et al. 2004, Rizzolatti and Sinigaglia 2008, Sinigaglia 2009). The general idea is that rudimentary forms of social cognition can be explained with the help of the mirror neuron system. The mirror neuron system (or, rather, systems, as there are many mirror neuron systems in the brain, but I will focus on the one in the parieto-frontal network) consists of bimodal neurons that get activated both when the agent performs an action, and when she perceives another agent performing this action (both in rhesus monkeys and in humans, Gallese et al. 1996, Umiltà et al. 2008, see Rizzolatti and Sinigaglia 2008 for a summary).

Importantly, the mirror neurons do not get activated when the perceived agent does not perform a goal-directed action but exhibits a mere (not goal-directed) bodily movement (Kakei et al. 2001, Umiltà et al. 2008). If the other agent is grasping a ball, the mirror neurons fire; if she is making a grasping movement without there being anything to be grasped, they do not. The general proposal then is that the mirror neuron system is capable of calculating the other agent's intention from her bodily movement. In this sense, it is capable of attributing a mental state (an intention or a "goal-state") to another agent.

There are some important worries about this general suggestion (not about the existence of mirror neurons, but about whether they can explain our cognitive engagements with others). I will mention only one (Jacob 2008, 2009, Jacob and Jeannerod 2003, see also Csibra 2007): it is not clear what serves as the input for the mirror neuron system and what we should consider the output. The standard story is that agent A observes the behavior of agent B and, on the basis of B's fine-grained movements, A's mirror neuron system computes B's intention. But this is not the only way of interpreting the empirical data. It is also a possibility that A has some kind of background representation of B's

intention and, on the basis of B's fine-grained movements and this background representation of B's intention, A's mirror neuron system predicts B's next move (see especially Jacob 2008).

Whether or not the explanation of our cognitive engagement with others with the help of mirror neurons is a viable one, it needs to be pointed out that it is very different from my approach. The main difference is this: mirror neurons get activated only if another agent is performing an action. But vicarious perception can happen when the other agent does not do anything at all. In the tiger example above, the other agent is standing still (while the tiger, unbeknownst to her, is attacking from behind), but we are nonetheless engaging with her very strongly. The mirror neuron hypothesis, even if it is unproblematic otherwise, cannot explain this.

Again, my proposal is not that we should replace all these accounts of social cognition (theory of mind, direct perception, mirror neurons) with a theory of vicarious perception. Social cognition is not a monolithic category: presumably, there are very different mental capacities involved in making sense of others. Some of these involve the attribution of mental states to others. Some others involve mirror neurons. My aim is to give a new account in addition to these, and argue that many rudimentary and (phylogenetically as well as ontogenetically) basic cases of social cognition are best explained by this new account. There are many ways of engaging with others, but vicarious perception is a very important and thus far ignored way of doing so.

6.3 Vicarious perception in the cognitive sciences

Let us return to the three grand questions about theory of mind: (i), (ii), and (iii). My claim is that if we replace the concept of theory of mind with vicarious perception, we get straightforward and nontrivial answers to the following new questions:

  (i) Are non-human animals capable of vicarious perception?
 (ii) How does vicarious perception develop in ontogeny?
(iii) What mental processes make vicarious perception possible in adult humans?

Let us take these questions in turn.

6.3.1 Vicarious perception in non-human primates

It has been severely debated in primatology, and cognitive science in general, whether chimpanzees have a "theory of mind": that is, whether they are capable of attributing mental states to other agents (Premack and Woodruff 1978, Heyes 1998, Gómes 1996, Tomasello and Call 1997, Tomasello et al. 2003, Povinelli and Vonk 2003, Andrews 2005, 2012, Byrne and Whiten 1988, Whiten 1996, Whiten and Byrne 1997).[2] There seems to be consensus on only one claim in this literature: namely, that on some interpretations of what theory of mind is chimpanzees do have theory of mind, whereas on others they don't (see Call and Tomasello 2008 for a summary).

The best candidate for primate theory of mind is "perspective taking": it has been argued that chimpanzees respond differentially to what they take the other agent to see (Hare et al. 2000, 2001, 2006, Tomasello et al. 2003, Call and Tomasello 2005, Tomasello and Call 2006, Kaminski et al. 2004, 2008, Melis et al. 2006, Tomasello and Call 1997). The key experiment here is one where a dominant and a subordinate chimpanzee spontaneously compete for food. One food item is visible to both of them, whereas the other one is only visible to the subordinate chimpanzee. It turns out that the subordinate chimpanzee vastly prefers to go for the food item that the dominant chimpanzee does not see. Although the results of this experiment have been disputed (see especially Karin-D'Arcy and Povinelli 2002), they have been successfully replicated with sufficient controls (Bräuer, J. et al. 2007).

A straightforward way of explaining this behavior would be to say that the subordinate takes into consideration what the dominant chimpanzee sees. In short, it attributes a perceptual state to the dominant chimpanzee.

Should we then conclude that chimpanzees are capable of attributing perceptual states to others? The problem is that some other experimental findings contradict this claim. Daniel Povinelli and his co-workers conducted a series of experiments that seem to demonstrate that chimpanzees do not attribute perceptual states to each other (Povinelli and Vonk

---

[2] There are also debates about the theory of mind capacities of other non-human animals, such as corvids; see, for example, Emery and Clayton 2004.

2003, 2004, Povinelli and Eddy 1996, Penn and Povinelli 2007, Reaux et al. 1999, see also Bulloch et al. 2008). The most decisive set of experiments is the following: the chimpanzee can ask for food from two experimenters, only one of whom seems to be able to see (say, because the other one has a bucket on her head). It turns out that chimpanzees ask for food from the two experimenters, regardless of which one appears to be a perceiver. These experiments seem to demonstrate that chimpanzees do not attribute even perceptual states to others. If they did, they would show a preference for asking for food from those agents who seem to be able to see them. But they do not show any such preference.

Hence, even the most plausible candidate for theory of mind—namely, the attribution of perceptual states to others—lacks conclusive support. We have reached an impasse (see also Sterelny 2003).

And here is the point where the concept of vicarious perception can be applied successfully. The Hare et al. experiments can be appropriately described as instances of vicarious perception. The subordinate sees the food item as affording the action of eating to the dominant, and she sees the other food item as not affording the action of eating to the dominant. She attributes very different other-centered—that is, dominant chimp-centered—action-properties to the two objects. And she, understandably, goes for the one that does not afford eating to the dominant chimp.

Note that many, maybe even most, experimental findings in favor of the existence of theory of mind in non-human primates work on the very same model, starting with Premack and Woodruff's original ones: the chimp in that experiment sees objects (a key, a tap, a switch, a box) as affording actions (opening the cage, rinsing the dirty floor, lighting the unlighted heater, reaching the bananas) to another agent (Premack and Woodruff 1978).

And the same is true of the alleged anecdotal evidence for theory of mind in non-human primates. One of the earliest such anecdotes is the following (Goodall 1971): a young chimpanzee noticed a piece of food on a tree, but did not get it and did not even look at it so long as there were other chimps around, but when they left she immediately jumped up and grabbed it (see also Byrne and Whiten 1988 and de Waal 1982 for similar anecdotes).

The Povinelli experiments are somewhat more difficult to handle. An important contrast here is the following set of experiments on rhesus

monkeys (Flombaum and Santos 2005).[3] The experimenters looked exactly the way they looked in the Povinelli experiments: one blind-folded, the other not blindfolded, for example. They even stood in a way very similar to the experimental setup in the Povinelli scenarios. The big difference was that the monkeys were in a competitive situation: the monkeys wanted to snatch the food from the experimenter, instead of asking him to give it to them.

Note that the Flombaum and Santos experiments fit the vicarious perception pattern perfectly: the monkey sees the food as affording an action (of taking it) for one of the experimenters, but not to the other. And as the monkey is afraid of the experimenter, she will only go for the food if it does not afford this action to the experimenter. This is a structurally very similar situation to the one in the Hare et al. experi-ments. The fact that the way the experimenters look is almost identical to the Povinelli experiments, however, offers us a neat way of comparing the alleged pro and the alleged con cases of primate theory of mind.

According to a popular and widely explored proposal, apes and monkeys tend to do much better at competitive mind-reading tasks than collaborative ones. The Hare et al. experiment is a competitive situation: the subordinate chimp is competing with the dominant one for food. The anecdotal evidence described in Goodall 1971 is also a competitive one, and so is the Flombaum and Santos experiment. The Povinelli experimental setup, in contrast, is a collaborative situation.

While this contrast has been made repeatedly in the literature (see Tomasello et al. 2005 for a summary), it is important to note that appealing to the difference between competitive and cooperative situ-ations only shifts the explanatory burden. If chimps do not attribute perceptual states in the Povinelli experiments because it is a cooperative situation, and they do so in the Hare et al. experiments because it is a competitive situation, this just raises an even more fundamental ques-tion: why should the competitive/cooperative distinction make a differ-ence in the attribution of perceptual states? Further, if chimps do have the capacity to attribute perceptual states (as evidenced by the competi-tive situations), what stops them from using this capacity in cooperative situations if it benefits them?

---

[3] The Flombaum and Santos experiment is about rhesus monkeys and not chimpanzees, but as rhesus monkeys are widely held to be less capable of theory of mind than chimpan-zees (see Cheney and Seyfarth 1990, Cheney et al. 1995), this asymmetry can be ignored.

One major advantage of the vicarious perception account is that it may help us to elucidate this distinction. First, I need to make it clear that we have no reason to assume that the chimp in the Povinelli experiments sees the food as affording an action to the experimenter. A much simpler and more parsimonious explanation would be to say that she sees the food as affording an action to herself, and the experimenter merely figures in this action as a means of achieving its goal. We can fully explain the chimp's behavior without talking about vicarious perception (a claim Povinelli would wholeheartedly agree with). In the Hare et al. and the Flombaum and Santos experiments, in contrast, we need to posit a vicarious perceptual process: the chimp needs to see the food as affording an action to the experimenter in order to respond differentially.

The big difference between the Povinelli setup and the Hare et al. setup, then, is the following. In both of these scenarios there are two potential actions: the food is perceived as affording one potential action to the dominant chimp/the experimenter, and it is perceived as affording another potential action to the (subordinate) chimp. In the Hare et al. scenario, these two actions are incompatible: it is either the dominant or the subordinate chimpanzee who gets to eat the food. In the Povinelli scenario, the actions are not incompatible: in fact, the experimenter's action is a way of performing the chimp's action. That is why there is no vicarious perception in the Povinelli scenario, while there is vicarious perception in the Hare et al. scenario.

To generalize from these two cases, in the case of competitive scenarios, the action the other agent (is about to) perform is incompatible with the action the subject (is about to) perform. In cooperative scenarios, there is no such incompatibility. Hence, in competitive situations, the subject needs to consider what action objects afford to the other agent—she can't just consider her own action. In the cooperative situation, she does not need to do so—she can just see the objects around her as affording actions to herself, and can see the other agent as a potential means to achieving her goal.[4] The real difference between competitive and cooperative situations is about actions—more precisely, about the

---

[4] Note that the claim is that the subject does not need to consider what action objects afford to the other agent. But the subject can consider this, which may be responsible for the discrepancies in replicating the original Povinelli experiments, see Bulloch et al. 2008.

compatibility/incompatibility of the actions afforded to the subject and to the other agent. The theoretical framework of vicarious perception can explain this difference. The theory of mind framework cannot.

I have argued that we can explain the seemingly conflicting results of the Hare et al. and the Povinelli experiments as instances of perceiving objects around the other agent as affording actions to this agent. Further, it seems that chimpanzees are capable of this (and maybe only this) way of engaging with other agents. Some of the controversy around the empirical findings about chimpanzee theory of mind is therefore terminological: if perceiving that objects around the other agent afford actions to this agent counts as having theory of mind, then chimpanzees have theory of mind. If not, they don't.

Another widely disputed experiment about social cognition in animals is about imitation in Japanese quail (Akins and Zentall 1998, see also Dorrance and Zentall 2001, Zentall 2003). In the most important experiment, "observer" quails watched other quails operating a simple mechanism: a lever that could be pushed down either by their beak or their feet. In half of the cases the other quail was rewarded with food, in the other half it was not. When the "observer" quail was allowed to operate this lever itself, only those observer quails that watched the other quails being rewarded imitated their behavior—the ones that watched the other quails going unrewarded did not. This experiment is supposed to demonstrate that the observer quail's behavior depends on whether the other quail received a reward.

It has been debated whether the observer quail's behavior can be explained by mere associative learning (in other words, without any appeal to social cognition)—there seem to be good reasons to believe that it cannot (Heyes 2001, Nehaniv and Dautenhahn 2002, Papineau and Heyes 2006). But it is also questionable that we should interpret the quail as making inferences on the basis of whether the other quail has been rewarded (or that we should explain the quail's behavior in terms of theory of mind) (Heyes 2005, Papineau and Heyes 2006).

If we accept the theoretical framework of vicarious perception, then we can easily explain the quail's behavior without over-intellectualizing it and without resorting to associative learning. The observer quail sees the lever as affording an action to the other quail. If the other quail is rewarded, then the action it perceives the lever as affording is very significant: it affords eating. If there is no reward, then the observer quail

perceives it as affording a much more boring action (of pushing the level with no apparent result). And when the observer quail is subsequently allowed to operate the lever, it perceives it as affording the action it has previously seen it as affording to the other quail. That is why those observer quails that have seen the other quail getting rewards are more likely to imitate its behavior.

6.3.2 *Vicarious perception and cognitive development*

One big debate in the theory of mind literature concerns developmental psychology: at what age do children acquire the ability to attribute mental states to others? The initial response was age 4 (Wimmer and Perner 1983, see also Wellman et al. 2001), which quickly went down to 2 (O'Neill 1996, Southgate et al. 2007), and to 1.5 years (Meltzoff 1995). More recently, there are more controversial proposals for evidence of social cognition in 15-month-olds (Onishi and Baillargeon 2005), 13.5-month-olds (Song et al. 2005), 13-month-olds (Surian et al. 2007), 12-month-olds (Gergely et al. 1995, Kuhlmeier et al. 2003, Csibra et al. 2003), 10-month-olds (Hamlin et al. 2007), 9-month olds (Csibra et al. 1999), and even 6.5-month-olds (Kamewari et al. 2005, Csibra 2008)—to mention only a few important milestones.[5] None of these proposals are uncontroversial—in each case there are suggestions for explanations of the displayed behavior without any reference to anything reminiscent of theory of mind (some important examples can be found in Perner and Ruffman 2005, Premack and Premack 1997).

It is important that not all of these experimenters are taking themselves to establish that infants attribute beliefs to others (although many do; see, for example, Surian et al. 2007): they loosen the criterion for theory of mind in a number of ways. Some talk of the attribution of goals (Csibra et al. 1999), of "perspective taking" (Tversky and Hard 2009), of the attribution of dispositions to act in a certain way (Hamlin et al. 2007, Song et al. 2005) or of "taking the intentional stance" (Gergely et al. 1995). It is not entirely clear how these proposals relate to one another and to the original question of the development of theory of mind.

---

[5] This list is by no means exhaustive—there are dozens of important studies about infant theory of mind published every year. Also, the age of 6 months is not the earliest age where the emergence of (something like) theory of mind is postulated. It has also been argued recently that it develops as early as at 3 months (see Luo 2011).

My suggestion is that all the relevant experimental findings that are supposed to demonstrate that 1-year-old and younger infants display the capacity to attribute mental states to others are in fact instances of vicarious perception. The question about the development of theory of mind may or may not have a straightforward answer, but we can use these experiments to give a fairly precise answer to the question about the development of a more rudimentary ability to engage with others, if we take this rudimentary form of social cognition to be vicarious perception. I argue that the capacity for vicarious perception develops somewhere between 6 and 9 months.

Let us take some of the most famous experimental findings about the early development of social cognition:

(a) Infants 13.5-months-old, who have watched an actor slide toy trucks on the floor, look at an actor who grasps a toy truck that is enclosed in a small frame longer than they look at an actor who grasps an identical toy truck that is not enclosed (and is therefore free to slide) (Song et al. 2005). The authors conclude that 13.5-month-olds attribute dispositions to act.

(b) One-year-olds, who have watched a circle go around an obstacle on its way to a larger circle, look at the small circle taking the detour when the obstacle is not there longer than they look at it going straight towards the larger circle without any detour (Gergely et al. 1995). The authors conclude that 1-year-old infants "take the intentional stance." The same is true of 9-month-olds (Csibra et al. 1999).

(c) Twelve-month-old infants prefer helpers to hinderers: if they see a triangle helping a circle up a slope, and a square trying to prevent the circle from climbing up the slope, they show preference for the triangle (Kuhlmeier et al. 2003). The authors' conclusion is that 1-year-olds evaluate others on the basis of their social behavior, and they attribute dispositions to others. The same is true of 10-month-old and 9-month-old infants, but not 6-month-olds (Hamlin et al. 2007).

(d) The general setup of experiment (b) was replicated with 6.5-month-olds, but the agent to whom the infants attributed goals was a human (not a small circle) (Kamewari et al. 2005). Later, it has been shown that the agent does not need to be a human (or even human-looking): it can be a box, as long as the route this box takes around the obstacle is varied (Csibra 2008).

A striking feature of these experiments is that they all seem to follow the same pattern, which is in fact the pattern of vicarious perception: the infant perceives an object as affording an action to another agent. More specifically, *the infant perceives* the toy truck (a), the obstacle or lack thereof (b), (d), and the triangle or the square (c) *as affording the action* of sliding (a), of going around it (b), (d), and of helping or hindering (c) *to* the actor (a), the circle (b), (c), and the box (d).

In other words, experimental findings (a)–(d) can be explained without any reference to the attribution of any mental state (be it belief, disposition, or goal). They can be explained with the help of a simple perceptual process: vicarious perception. The evidence for social cognition in infants younger than 1 year is in fact evidence for vicarious perception in these infants. And this evidence tells us that vicarious perception emerges between months 6 and 9.

Take (c) as an example. The data is that the infants show preference for the triangle that helps the circle up the slope over the square that is trying to prevent the circle from climbing up the slope. We can explain this by describing the infant as having a non-perceptual mental state of attributing a disposition or maybe even virtue/vice to the triangle and the square. But we also have a much simpler explanatory scheme, with reference to the infant's perceptual states alone: the infant does not attribute any mental state (or disposition) to anyone, she merely perceives the triangle as affording a certain action to the circle, whereas she perceives the square as affording another action to the circle (these geometrical figures are all taken to be agents—see Scholl and Tremoullet 2000). On the basis of these vicarious perceptual states, she forms an understandable preference for the triangle. Examples (a), (b), and (d) can be analyzed in the same way.

To sum up, the experiments for early social cognition in developmental psychology say little about the attribution of mental states. They do, however, give us a firm understanding of the emergence of vicarious perception in infancy.

6.3.3 Vicarious perception in adult humans

The third big debate about the concept of theory of mind is not about preverbal infants or non-human primates, but about adult human beings. We attribute beliefs and desires to each other all the time. The question is: how do we do it? What are the mental processes that make the cognitive engagement of adult humans with others possible?

There seem to be two candidates, the "simulation-theory" and the "theory-theory." The latter claims that we are equipped with a theory whose domain of application is constituted by other agents' mental states. On this view, attributing a mental state to someone else is a case of applying a psychological theory. This is the theory-theory view (Nichols and Stich 2003, Botterill 1996).

According to the alternative account, we have the capacity to simulate other people's mental states; that is, we are able to put ourselves in other people's shoes and go through, in imagination, the mental states we would go through were we really in the other person's circumstances. The end result of such a process, namely the mental state in which the simulator finds herself, can now serve as a guide to what mental state the simulated person is in. This is the simulation view (Gordon 1995a, 1995b, Heal 1995, Stone and Davies 1996, Goldman 1992, 2000, 2006).

While there are some important disagreements about what should be meant by simulation, a good approximation is the following: an agent A imagines herself in B's circumstances, gets a grip on what she, A, would do (see, feel, think, and so on), and concludes that this is what B would also do (see, feel, think, and so on) in these circumstances. As Gregory Currie writes: "I imagine myself to be in the other person's position, [. . .] I simply note that I formed, in imagination, a certain belief, desire or decision, then attribute it to the other" (Currie 1995b, pp. 144–5). But it is important to note that this "putting or imagining oneself in the other person's position" is not necessarily conscious or explicit (see Goldman 2006 and Gallese and Goldman 1998, Keysers and Gazzola 2007, Gallese et al. 2004).

The debate about the theory of mind of human adults has been revolving about the respective merits, and especially demerits, of these two accounts. My suggestion is that this is a false dichotomy: we do not have to choose between simulation-theory and theory-theory (or find a hybrid account that would combine the two; see, for example, Heal 1998, Stone and Davies 1996, Perner 1996). Vicarious perception is a way of engaging cognitively with others that is very different from (and much simpler than) both simulation and the application of a theory.

The idea that some of our simplest and most basic capacities to engage with others is neither simulation nor the application of a theory, but something entirely different, is not new (Zahavi 2008, Gallagher 2005, 2007a, 2007b, Hutto 2007, 2008). My account differs from these recent

ways of resisting the dichotomy of simulation-theory and theory-theory in at least two ways. First, I do not claim, like some of these authors, that we directly perceive others' emotions and mental states (Zahavi 2008, Gallagher 2008, see also Section 6.2 for the distinction between this view and mine). Second, I do not claim that a species of narrative understanding is involved in understanding others (Hutto 2008).

We can engage cognitively with others in a variety of ways. Two of these have received a lot of attention: simulation and the application of a theory. But there are more. In short, we should not think of the question about the mental processes that make social cognition possible to be the question of simulation versus theory. We should open up this debate to include other forms of engaging with others, like vicarious perception.

One important empirical application of the debate about the mental processes that make social cognition possible concerns the "Social Simon Effect." In the standard Simon task, the participant carries out a spatially defined response to a non-spatially defined stimulus, but the location of this (non-spatially defined) stimulus influences the response time: responses are faster if stimulus location and response location correspond. For example, if the subject sees a triangle she is supposed to push a button on her right, and if she sees a square she has to push a button on her left. When the triangle appears in the right-hand side of her visual field (or if it appears together with some marker that emphasizes that side of her visual field), her reaction is faster than it is when it appears on the left. This is the standard Simon effect (Simon 1990). The Social Simon Effect replicates this result in a scenario where the two different responses are carried out by two different people: if there is another agent on my left pushing a button when seeing a triangle, my reaction to the square is faster when it appears on the right than it is when it appears on the left. This difference disappears if there is no-one on my left.

The original interpretation of the Social Simon Effect was that it demonstrates how we have "action co-representations" when we perform joint actions with others: a representation of both one's own action and the other person's action (Sebanz et al. 2003, 2006): "if co-actors represent each other's actions, an irrelevant spatial cue referring to the other's action should activate the representation of the other's action and create a response conflict" (Sebanz et al. 2005a, p. 1234). It is easy to see that having an "action co-representation" entails attributing a mental state to the other agent: it entails "theory of mind."

But this is not the only interpretation. It has also been suggested that the reason for the Social Simon Effect is that the other agent provides a spatial reference frame—the other person's mind does not play any role in creating the effect, she is relevant only for helping us to localize the stimulus and the response in space (Guagnanoa et al. 2010, Dolk et al. 2011). On this interpretation, the Social Simon Effect is not social at all— it does not involve any form of social cognition.

Neither of these interpretations is unproblematic. The problem with the first, "action co-representation," interpretation is that the Social Simon Effect is also present when the subject is a patient with autism spectrum disorder (Sebanz et al. 2005b). But it is widely held that autism can (at least partly) be explained in terms of the subjects' deficiency of "theory of mind" (Baron-Cohen 1995, Senju et al. 2009). But then how is it possible that they are capable of forming "action co-representations" (see also Humphreys and Bedford 2011)? Also, it turns out that the further away the agents sit from each other, the weaker the effect gets (Guagnanoa et al. 2010). It is not at all clear why this would make a difference if the effect is to be explained by a version of "theory of mind."

The problem with the second, "spatial reference frame," interpretation is that the Social Simon Effect depends on the actor's bad mood (Kuhbandner et al. 2010), and, importantly, her negative relationship to the other actor (Hommel et al. 2009). Further, if the agent believes that the co-actor is a computer, the effect disappears (Tsai et al. 2008). These findings indicate that there must be something "social" in the Social Simon Effect.

The concept of vicarious perception can help us to resolve this debate. The Social Simon Effect can be interpreted as a manifestation not of "theory of mind," but of vicarious perception. The effect is present because the actor is aware of the action the stimulus on the left-hand side affords to her co-actor. The actor sees the stimulus on the left-hand side as affording an action not to herself, but to her co-actor. The actor attributes other-centered—that is, co-actor-centered—action-properties to the stimulus.

This does not entail attributing any mental states to the co-actor— which explains why the effect is still present in the case of autism spectrum disorder patients. But it does involve social cognition: namely, vicarious perception—which explains why the effect is sensitive to the agent's mood, to the relationship between the agents, and to whether the

agent thinks that the co-actor is a computer. As in the case of non-human primates and pre-verbal infants, explaining the Social Simon Effect by appealing to vicarious perception is a novel, third way between explaining the effect with the help of (some version of) "theory of mind" and explaining it as not involving any social cognition whatsoever.

I argued that we should shift the emphasis in one of the most import-ant contemporary debates about social cognition in psychology and the cognitive sciences from theory of mind to vicarious perception. Should we as a result retire the concept of theory of mind? I don't think so. We adult humans do often attribute beliefs and desires to others. I tried to point out that this does not happen as often as we may think, since there are other, simpler (perceptual) ways of engaging cognitively with others, but it does happen. Theory of mind may not be the holy grail of understanding social cognition, but it labels a real and important phenomenon.

6.4 Vicarious emotions

There are many ways of engaging emotionally with other people, and there are many words philosophers and non-philosophers use to label them: empathy, sympathy, fellow-feeling, *Entfühlung*, etc. The aim of this section is not to give an analysis of any of these terms, but to highlight the importance of a way of emotionally engaging with others that has received much less attention.

This unexplored way of emotionally engaging with other people is a variant of vicarious perception. When we perceive something vicari-ously, we see it as affording an action to someone else. But sometimes we perceive entities around another person not as affording an action, but as *emotionally relevant* to this person. I argue that this way of emotionally engaging with another person is importantly different from the existing accounts, and may also be more basic both evolution-arily and developmentally.

A couple of terminological remarks: I use the term "emotional engage-ment with others" as an umbrella term to refer to the wide variety of affective phenomena. Very different affective states fall under this umbrella term: empathy, sympathy, fellow-feeling, and more. I do not aim to give an exhaustive account of every kind of emotional engagement

with others—I don't think this would be a feasible task. I do not aim to give an account of empathy or of sympathy either—I am not sure that either of these concepts refer to emotional natural kinds. I will use the term "vicarious emotional engagement" to refer to the form of emotional engagement with others in which we perceive entities around another person as emotionally relevant to that person.

6.4.1 Vicarious emotional engagement

When I look at a cockroach crawling up my leg, I feel disgust. When I look at the neighbor's huge pit bull running towards me, I am afraid. One aspect of what goes on in my mind during these moments is that I represent the cockroach as disgusting or the dog as scary. In general, when we have an emotion directed at an object, there are some properties that we need to represent this object as having in order to have this emotion. I call these properties "emotional-relevance properties."

The term "emotional-relevance property" is a technical term: it is a convenient label to use for properties like being scary (in the case of the emotion of fear) or being disgusting (in the case of the emotion of disgust). When we have an emotion directed at an object (say, when we are disgusted by a cockroach), we may represent it as having all kinds of properties: shape, size, color, etc. But we can have this emotion without representing this object as having some of these properties. We could not have the emotion of being disgusted by x without representing it as having the property of being disgusting: without representing it as having an emotional-relevance property.

It is important to emphasize that this is not in any way a new account of emotions, nor does this make us side with one theory of emotion or another. This is just a way of talking about emotions that allows us to focus on the attributed properties.

Importantly, this way of talking about emotions does not imply that all emotions are necessarily representational. There is a big debate in the philosophy of emotion literature about whether every emotion is necessarily directed at certain events or entities: whether they are necessarily *about* something (which would provide a contrast with moods, since moods are not directed at certain events in the same way). There is also a debate about whether the representational properties fully capture what an emotion is (or maybe emotions need to (also) have some non-representational aspect) (see Hume 1739/1978, James 1884, Oatley and

Johnson-Laird 1987, Pitcher 1965). I want to bypass both of these debates. Maybe there are genuine emotions that do not represent any event or entity. But in the case of those emotions that *do* represent events or entities, these emotions involve (but may or may not be constituted by) the representation of these entities as having emotional-relevance properties.

Further, I also want to bypass another big debate about the nature of emotions; namely, about whether there are unconscious emotions (see, for example, Winkielman and Berridge 2004, Hatzimoysis 2007). That is why I talk about the *representation* of an entity or event as having emotional-relevance properties. If emotions are necessarily conscious (something I am not at all committed to), we can rephrase "representing x as having an emotional-relevance property" as "experiencing x as having an emotional-relevance property." As I want to leave open the possibility of unconscious emotions, I will use the neutral concept of representation in what follows.

Emotional-relevance properties, like thick action-properties, are relational, and their attribution is a highly contextual affair: the same subject may attribute very different emotional-relevance properties to the very same entity depending on her mood, the other entities surrounding her, and other contextual factors. For the ease of exposition, I will use the phrases "A represents x as having an emotional-relevance property" and "A represents x as emotionally relevant to her" interchangeably.

So far, I have been describing an agent's representation of something as having self-centered emotional-relevance properties: as being emotionally relevant to herself. But we also often represent entities as having other-centered emotional-relevance properties: as being emotionally relevant to other agents. This is the form of emotional engagement I call "vicarious emotional engagement."

Suppose that you see your friend, Bill, sitting on a bench and, unbeknownst to him, a rat is approaching him from behind and starts sniffing at his shoes. You know that Bill hates rats. This is a form of emotional engagement, but it is unclear how the traditional accounts of empathy or other forms of emotional engagement would analyze this scenario. My way of describing what happens is that you represent the rat as emotionally relevant, not to yourself, but to Bill. You represent it as having a Bill-centered emotional-relevance property. You don't represent it as emotionally relevant to you because it is far away from you and,

let us suppose, you have no problem with rats. But you know that Bill finds them disgusting and, as a result, you represent it as emotionally relevant to him. Vicarious emotional engagement is a very simple, visceral, quasi-automatic, and, arguably, perceptual process. This, I claim, is a thus far underrated way of engaging with others emotionally.

Take another example: you go to a party with your friend, Fred, who has just had a messy divorce from his wife, Jane. Entering the party first, you see Jane there kissing another man. Again, what happens here is that you represent the event of Jane kissing another man as emotionally relevant, not to yourself, but to Fred. You attribute Fred-centered and not self-centered emotional-relevance properties to this event. That is, you don't represent this event as emotionally relevant to yourself—you've never had any strong feelings for or against Jane. But you know that Fred cares about her a great deal and, as a result, you represent this event as emotionally relevant to him. This example illustrates that one's vicarious emotional engagement can be (but doesn't have to be) sensitive to one's higher-order thoughts and beliefs, such as my beliefs about Fred and Jane's past (see de Sousa 1987).

Note the structural similarity with vicarious perception. Vicarious perception is the perceptual attribution of other-centered thick action-properties to an object. Vicarious emotional engagement is the perceptual attribution of other-centered emotional-relevance properties to an object. Further, vicarious emotional engagement does not entail the attribution of an emotional state to the other agent: neither Bill nor Fred attributes emotional-relevance properties. Fred is still coming up the stairs and hasn't seen what I have seen. And Bill is not aware of the rat. So my vicarious emotional engagement with them does not entail the attribution of an emotional state to them (if it did, this would be a misattribution). Vicarious emotional engagement can be accompanied, and as a result colored, by the attribution of emotional states (for example, by my attribution of Fred's general emotional attitude towards Jane), but it does not have to be. Like vicarious perception, vicarious emotional engagement is not necessarily accompanied by the attribution of mental (in this case, emotional) states.

But the relation between vicarious perception and vicarious emotional engagement is more than a mere structural parallel. As we have seen, vicarious perception, like any attribution of thick action-properties, can be a highly emotional affair (and it can also be influenced by our

emotions, see Morgado et al. 2011): when we see something as affording an action to our friend, this tends to be accompanied by affective phenomenology—as in the case of the tiger example above. Further, many (basic) emotions are genuinely action-oriented (see Prinz 2004 for a summary): seeing something as disgusting is partly constituted by our propensity to avoid it; seeing something as scary is partly constituted by our propensity to escape, etc.[6] In short, there is significant overlap between vicarious perception and vicarious emotional engagement.

One may wonder whether the attribution of emotional-relevance properties, like that of action-properties, is a perceptual process. Without committing myself to the claim that it is, it is worth pointing out that there may be good reasons to think so. If emotional-relevance properties were attributed non-perceptually, this would make it difficult to explain the well-documented belief-insensitivity of our emotions (Greenspan 1988, see also Zajonc 1984 and Prinz 2004, p. 35). Maybe an argument, structurally similar to the one I gave in Chapter 2 about the perceptual attribution of action-properties, could be constructed, but I will not attempt to do this here.

Further, there are some empirical findings that point in the direction of the attribution of emotional-relevance properties as being perceptual. The representation of emotional-relevance properties can be and has been examined experimentally. The physiological changes in our body are different depending on whether we are shown pictures of emotionally salient events (like romantic couples kissing, or mutilation) or pictures of emotionally neutral events. Most importantly, there are positive voltage changes in event-related brain potentials when the pictures shown are emotionally salient, but not when the pictures are emotionally neutral (Cuthbert et al. 2000, Schupp et al. 2000, Keil et al. 2002).

Importantly, the most accurate measure of whether the agent represents the contents of the picture she is looking at as having emotional-relevance properties is the way her perceptual attention is exercised. If she is looking at emotionally salient pictures, her blinking reflex to a distracting noise or flash is longer and slower than it is when she is looking at emotionally neutral images (Schupp et al. 2004). But if the representation of emotional-relevance properties can be measured by

---

[6] For the distinction between basic and non-basic emotions, see Damasio 1994, Griffiths 1997, 2004, but see also Ortony and Turner 1990 and Clark 2010 for criticisms.

measuring the engagement of one's visual attention, then a strong case can be made for the perceptual representation of emotional-relevance properties.

I do not think that this empirical data provides a knock-down argument for the view that emotional-relevance properties are perceptually represented. There may be ways of accommodating these empirical findings while holding on to the claim that emotional-relevance properties are not perceptually represented: for example, if one maintains that the perceptual representation of shape and color gives rise to a non-perceptual representation of emotional-relevance properties, which in turn has some kind of top-down influence on visual attention. But taking this route would imply the entirely ad-hoc postulation of explanatorily superfluous non-perceptual representations, and it would also imply endorsing a revisionary view of visual attention (see Nanay 2010a).

As already stated, I will not assume here that the attribution of emotional-relevance properties (including other-centered emotional-relevance properties) is a perceptual process. But it is a relatively simple process that does not require higher-order thoughts. After pointing out that vicarious emotional engagement is different from, and simpler than, other forms of emotional engagement with others, I argue that if non-human animals are capable of engaging with others emotionally at all, this is likely to be of the form of "vicarious emotional engagement."

6.4.2 Vicarious emotional engagement versus simulation

The aim of this subsection is product differentiation: to point out how my account of vicarious emotional engagement differs from existing accounts of engaging emotionally with others. There are many theories of our emotional engagement with others. I will mainly focus on what I take to be the dominant, simulationist account.

It is widely assumed that empathy is a form of simulation (see Adams 2001, Gordon 1995a, 1995b, Ravenscroft 1998, see also Currie and Ravenscroft 2002 for a summary, but see also Deonna 2006, 2007). As we have seen, simulation is the following mental process: an agent A imagines herself in B's circumstances, gets a grip on what she, A, would do (see, feel, think, and so on) and concludes that this is what B would also do (see, feel, think, and so on) in these circumstances (Currie 1995b, pp. 144–5).

The first thing to note about the simulationist account of empathy is that, on the face of it, it has little to do with emotion. In fact, the concept of simulation in the philosophy of mind was originally used in the debate about how we attribute mental states to others (Davies 1994; Carruthers and Smith 1996). The question is then what the simulationist attribution of mental states has to do with empathy. Different simulation theorists answer this question differently. One possible response is that just as we attribute beliefs and desires to others, we also attribute emotional states to others. Another possible, somewhat unorthodox, response is that the concept of empathy does not have to be restricted to emotional engagement (Gallese 2005, 2007). I put these worries about the simulationist theories of empathy aside (see also Zahavi 2008 for further worries), and I also bracket some more general worries about the simulationist framework in general (Gallagher 2001, 2005, 2007a, 2007b, Hutto and Ratcliffe 2007, Hutto 2004, 2007, 2008, Ratcliffe 2007, Hutto and Myin 2013).

There are many differences between my account of vicarious emotional engagements and the simulationist account of empathy. The most important of these is that while vicarious emotional engagement can be described in very simple (arguably) perceptual or quasi-perceptual terms, simulation requires some version of imagining from the inside and mental state attribution. As we have seen, vicarious emotional engagement does not even presuppose the attribution of an emotional state to another agent. Not only does the simulationist account of empathy presuppose such mental state attribution, it even supplements it with a complex imaginative process. In other words, my account describes a much simpler, less cognitively demanding way of emotionally engaging with others than the simulationist theories of empathy. I am not denying that some ways of engaging emotionally with others can be captured by the simulationist account of empathy, but some more rudimentary cases of emotional engagement with others cannot (I agree here with Gallagher and Meltzoff's 1996 objections to simulationism, without endorsing their positive alternative).

According to most simulationists, the imaginative episode necessary for empathy amounts to a version of "imagining from the inside" (Darwall 1998, Gordon 1995a, Walton 1990, p. 255, Currie 1995b, p. 153, Wollheim 1974, p. 187, Wollheim 1987, pp. 103, 129, Neill 1996, Smith 1997, Gaut 1998, cf. Feagin 1996, pp. 113–42). The proposal is that when I empathize with X, I imagine X from the inside.

As we have seen, the most promising version of the simulationist account is one where imagining X from the inside means imagining oneself being in X's situation, an account that can be traced back to Adam Smith's account of sympathy (Smith 1759/1976, see also Gaut 1999, Goldman 2006, Steuber 2006). Hence, if we want to explicate what this concept entails, we need to explicate what is meant by "X's situation." Depending on the way we interpret this notion, we end up with very different accounts of imagining from the inside and of empathy.

The first thing to notice is that we should not restrict X's situation to X's physical situation. X's situation should also include facts about what X knows. Suppose X is attacked by someone. The experience of imagining myself in X's situation will depend on whether X has a gun in her pocket, as this is an important element of X's physical situation. Similarly, the experience of imagining myself in X's situation will also depend on whether X knows something about the attacker that could be a means of defending herself (say, by blackmailing). And this is not an element in X's physical, but epistemic situation. In short, sometimes X's situation should include facts about X's psychology or about what X knows.

But sometimes it shouldn't. Many uncontroversial examples of empathy seem to entail that we actively *ignore* the psychological elements of the other person's situation. When we empathize with lunatics or with the dead, we need to ignore the psychological elements in their situation, a point also first noted by Adam Smith (Smith 1759/1976, I. i.1.11, see also Darwall 1998, Nanay 2010e). Most of the examples I gave for vicarious emotional engagement in Section 6.4.1 are scenarios where there is an important epistemic difference between the two agents: I am aware of the cockroach on your left leg but you are not.

In these cases, if it is possible at all to feel empathy towards others, one must imagine oneself in her situation, not as actually presented to her, but as presented to her were she to know what I know about her situation. The same point applies in the case of our emotional engagement with lunatics and the dead. Thus, when one imagines oneself in someone else's situation, one may need to abstract away from the psychological elements in that person's situation. We can and do empathize with people even if we know significantly more (or significantly less) about their situation than they themselves do. And, oddly, our emotional reaction in these cases seems to be much stronger than it would be if that person were aware of the same things we were aware of.

It is important that in these cases of empathy, we need to be aware of the other person's situation as it is presented to her, and nonetheless abstract away from it. If we are not aware of the other person's situation at all, then we tend not to be able to empathize with her: yet another point that was anticipated by Adam Smith (Smith 1759/1976, I.i.1.8). But this means that in cases where I know something about the situation of the person I am empathizing with that she herself does not know (and I know this to be the case), the empathizer needs to be able to have a representation of the other person's situation as it is presented to her and abstract away from this, to take into consideration a piece of knowledge that the empathizer is aware of but the other person is not, and represent this piece of knowledge as something the other person is not aware of.

This way of engaging emotionally with others then requires an extremely intricate ensemble of representations, most of which are detached from the sensory stimulus, and an equally complex exercise of one's imagination. This account then will not be a very good candidate for describing the emotional engagements of small children and of non-human primates (if we can talk about emotional engagement in the case of non-human primates, see Section 6.4.3). Simulationist accounts of empathy may be the right way of describing some, fairly complex, episodes of emotional engagements with others, but it is unlikely to be the right account of our simple, quasi-perceptual, quasi-automatic emotional engagements.

I said that the simulationist account of empathy is the mainstream account of our emotional engagement with others. But it is not the only one. I need to briefly contrast my account with some other (off-mainstream) accounts of emotional engagement that also appeal to our simple quasi-perceptual states (the following contrasts will be the structural equivalents of the ones I drew in Section 6.2).

The first such approach emphasizes the possibility of seeing someone else's emotion *directly*, on their face (McNeill 2012, Smith 2010, Goldie 1999, 2000, 2002, 2007, Dretske 1973, Cassam 2007). The general upshot is that by looking at X's face, I directly see her emotion. The success of this approach is intertwined with the success of its underlying assumptions in philosophy of perception, especially with regards to what is referred to as the "thickness" of perceptual experience. Without assessing its plausibility, it is important to point out that the phenomenon this approach describes is very different from vicarious emotional

engagement. In the case of vicarious emotional engagement, seeing the face of the other agent is not necessary for emotionally engaging with her. I can see the cockroach that is climbing up your leg as emotionally relevant to you even if I do not see your facial expression. Engaging emotionally with others on the basis of their facial expressions may very well be an important phenomenon, but not one I discuss in this chapter.

The other important (but still off-mainstream) approach uses empirical findings about the mirror neuron system in order to explain our engagement with others (Gallese 2005, 2007). Sometimes this approach is presented in conjunction with a version of simulationism, where the mirror neuron system is supposed to underlie our capacity to simulate others (Gallese and Goldman 1998). Again, leaving the plausibility of this approach aside (but see Debes 2010), it is important to point out that our mirror neurons fire when we see others perform a goal-directed action. What I call vicarious emotional engagement can and does happen even in the absence of seeing others perform goal-directed actions. When I see the cockroach climbing up your leg, you may be sitting on a bench doing nothing. Nonetheless, my vicarious emotional engagement towards you will be very strong.

6.4.3 Emotional engagement in animals

It is a controversial question as to whether we can talk about empathy or other forms of emotional engagement with others in the case of non-human animals. Although there is a lot of discussion of empathy in non-human primates (see Preston and de Waal 2002 and O'Connell 1995 for summaries), most of the discussion is about primates having some kind of behavioral reaction to the misfortunes of another primate. This behavioral reaction ranges from running away or hiding to trying to help. Besides primates, even albino rats and pigeons display similar behavior (Rice 1964, Watanabe and Ono 1986).

The problem is that it is not clear what would justify describing the animal as empathizing or even as engaging emotionally with the other. These experiments may show that animals are aware of the suffering of other animals, but, strictly speaking, they say nothing about whether this awareness gives rise to any kind of emotional engagement with the other animal (this holds even for the very suggestive anecdotal evidence in Ingmanson 2002, O'Connell 1995, and de Waal 2009, where the behavior

displayed by the animal uncannily resembles our own behavior—putting one hand on the other's shoulder, looking into her eyes, etc.).

As a result, some primatologists are skeptical about the very idea of animal empathy (see, for example, Byrne 1995, Cheney and Seyfarth 1990), and dismiss it as a clear case of anthropomorphizing animals. I argue that if non-human primates are capable of engaging emotionally with others, this is likely to take the form of vicarious emotional engagement.

Take the following experiment, which I take to be one of the strongest cases for emotional engagement in animals (Parr 2001). Chimpanzees who are shown short clips of other chimpanzees being injected with needles react with a sharp decrease of peripheral skin temperature (which indicates great negative arousal). There is no decrease of their peripheral skin temperature if the chimpanzee in the clip is shown in an emotionally neutral situation (see also Parr and Hopkins 2000, Winslow et al. 2002, Parr et al. 2002).

First, why should we think that this is an instance of emotional engagement with others? More precisely, couldn't we interpret these findings as the chimp's (conditioned) reaction to the sight of the needle? In other words, couldn't we say that the chimp's emotional reaction has little to do with the pain the other chimp suffers, and more to do with the sight of the needle that is potentially painful for the chimp herself?

There is a control experiment where the chimps are shown the needles only. Although these needles are more visible (and look much bigger), the result is a (slightly) less dramatic decrease in peripheral skin temperature. This makes it unlikely that their reaction is merely a conditioned reflex to the unpleasant sight of the needle. What explains the strength of their emotional reaction is that it is another chimpanzee who is being injected. I need to be clear that I do not take this experiment to be watertight proof that chimpanzees engage emotionally with others. I aim to make a much more modest conditional claim: if chimpanzees are capable of emotional engagement with others, this takes the form of vicarious emotional engagement.

Now the question is: how can we describe the results of this experiment? As the injected chimp's face is not always visible in these clips, the "seeing emotions in the other's face" explanation will not work. Nor will the mirror neuron explanation, as the injected chimp is not performing any action.

The simulationist account, in contrast, can explain these reactions, but it would give us a highly complicated picture of the chimpanzee's mental processes. Remember that according to the simulationist account of empathy, A imagines herself to be in B's situation, where B's situation is represented in such a way that it may include information that is available to A, but not to B. So if we want to describe the mental processes of the chimpanzee in this experiment in a simulationist way, we need to say that she imagines herself in the injected chimp's situation (an imaginative episode we have no evidence that non-human primates are capable of), and represent the injected chimp's situation in a way that would combine information that is available to the injected chimp and information that is not (a kind of decoupled representation we have little evidence that non-human primates are capable of either). In short, the simulationist explanation of this experiment, although possible, would amount to a serious over-intellectualization of the chimpanzee's mind.

I suggest that the chimp's emotional reaction can be described in a much simpler way: she attributes an other-centered emotional-relevance property to the needle; she sees the needle as emotionally relevant to the other chimp. It is important that the chimpanzees do not see the needle as emotionally relevant to themselves: otherwise we would get a stronger emotional reaction in the control experiment, where the needle is fully in view. They see it as emotionally relevant to another agent. The only mental processes that are needed to explain this emotional reaction are the ones that attribute other-centered emotional-relevance properties to a familiar object.

In short, if chimpanzees are capable of engaging emotionally with others (something I do not take this experiment to provide foolproof evidence for), this engagement is likely to be of the form I described as vicarious emotional engagement: seeing objects as emotionally relevant to another agent.

In conclusion, I outlined and defended a new account of a very simple way of engaging emotionally with others: we see entities and events around the other person as emotionally relevant to her. My aim was not to propose this account as a replacement for existing accounts of our emotional engagement with others, but rather in addition to them. There are many ways of engaging emotionally with others—my aim was to describe and analyze a so far unexplored way of doing so: seeing things as emotionally relevant to others.

6.5 Vicarious perception: the first step towards decoupled representation

The first four chapters of this book identified a kind of representation that plays an important role in the mental life of both humans and non-human animals: pragmatic representations. In Chapters 5 and 6 I examined two important variants of pragmatic representations— pragmatic mental imagery and vicarious perception—and pointed out that they play an important role in various human (maybe uniquely human) mental processes. But as pragmatic mental imagery and vicarious perception also attribute action-properties (or thick action-properties), all the mental phenomena this book has analyzed concern various ways of attributing various action-properties.

But we humans can and often do attribute properties that are not action-properties. In this last section I want to consider, very tentatively, how this arguably uniquely human capacity came about. In other words, I want to explore—again, very tentatively—another potential application of the idea of vicarious perception on some uniquely human capacities: on the origins of our capacity to have what are sometimes referred to as "detached" or "objective" mental representations (Sterelny 2003, Burge 2010).

The concept of detached representation has been used in a number of different senses (that are not always clearly differentiated). Sometimes it is used as a synonym for non-perceptual representation: representation that is detached from our sensory stimulation (Perner 1991, Leslie 1987). I did discuss, very briefly, the origins of representations that are detached in this sense when I talked about the evolutionary advantages of pragmatic mental imagery in Chapter 5. But this is not the sense of "detached" I am interested in here. Some of our representations are genuinely egotistic: they are about how the world is for us, not how the world is per se. Pragmatic representations are the paradigmatic examples: they represent only what is relevant to the performance of our actions. They are self-centered, egotistic representations.

The question is: how did we acquire the capacity to have representations that are less self-centered and less egotistic? How did we acquire the capacity to have representations that are detached from our own self and self-interest? As John Campbell says, "there is no reason why evolution should have fitted out animals with a capacity to represent particulars

that goes beyond their needs" (Campbell 2011, p. 282). What then made us acquire the capacity to represent particulars in a way that goes beyond our needs?

The very question of the origins of our objective representations is discussed recently at great length by Tyler Burge (Burge 2010). He criticizes the most influential contemporary accounts of what the capacity for objective representation presupposes (by Evans, Strawson, Quine and Davidson) for over-intellectualizing the mind. My approach is clearly similar to Burge's in this respect, but my positive account is very different from his.[7]

To use the terminology of this book, the question can be summarized in the following manner. Animals attribute action-properties to objects. This capacity makes a lot of evolutionary sense: the attribution of action-properties is by definition a survival-enhancing capacity. But why did we acquire the ability to attribute properties that are not action-properties? Why did we acquire representations that are detached not from sense perception, but from action?

My very sketchy and tentative answer is that we did so by means of the mediation of vicarious perception. Vicarious perception is not an "objective" or "detached" representation. It still consists of the attribution of action-properties. However, these action-properties are not self-centered, but other-centered. Vicarious perception still does not represent the way the world is per se. It represents how the world is *for another person*. But, and this is the crucial step, by representing both how the world is for us (pragmatic representation) and how the world is for someone else (vicarious perception), we are in the position to become aware of the difference between the other agent's perspective and our own, which in turn can lead to our representing the world as it is, and not as seen from one fixed perspective or another: as objective states of affairs.

This way of bringing in the social dimensions to explain the origins of objective, detached representations is, of course, not new. Its main proponent within philosophy was Donald Davidson (the same idea has become very influential in cognitive science; see, for example,

---

[7] Burge argues that perception itself can give us objective representations, which is certainly true under some interpretations of the concept of objectivity, but it is not clear to me that it is also true under the interpretation that I have in mind here—in the sense of being detached from the agent's actions.

Tomasello 1999), whose concept of triangulation was an attempt to capture the intersubjective origins of our objective, detached representations. As he says:

> Our sense of objectivity is the consequence of [a] sort of triangulation, one that requires two creatures. Each interacts with an object, but what gives each the concept of the way things are objectively is the base line formed between the creatures by language. (Davidson 1982, p. 327)

I agree with the first half of this quote: our sense of objectivity is the consequence of a sort of triangulation, and one that requires two creatures. One perceives the object as affording an action to the other. This already allows for the awareness of the difference between the other agent's perspective and one's own, which is a good starting point for representing the world not as seen from one fixed perspective or another, but as things are, regardless of perspective.

But according to my account, this form of triangulation does not require language (this is an idea Davidson himself briefly flirted with at the end of his life (personal communication), from Spring 2000 to Spring 2003, see also Davidson 2001, p. 128). In fact, all it requires is a perceptual process of seeing something as affording an action to someone else.

Thus, here is the (again, sketchy and tentative) picture that I take to be different both from Davidson's and from Burge's. First, we all had only self-centered pragmatic representations. Then some of us (not only humans) acquired the capacity of forming other-centered vicarious perception. And while this mental process is still nothing more than the attribution of action-properties, it can pave the way towards the acquisition of the capacity to form decentered (detached, objective) representations. Importantly, these "bird eye view" (Tomasello et al. 2005) representations or "naked intentions" (Jeannerod and Pacherie 2004) are later developments than vicarious perception. Nonetheless, vicarious perception may have been the stepping stone for the emergence of objective thought.

References

Adams, F. (2001) Empathy, neural imaging and the theory versus simulation debate, *Mind & Language*, 16, 368–92.

Aglioti, S., DeSouza, J. F. X., and Goodale, M. A. (1995) Size-contrast illusions deceive the eye but not the hand. *Current Biology*, 5, 679–85.

Akins, C. K. and Zentall, T. R. (1998) Imitation in Japanese quail: The role of reinforcement of demonstrator responding. *Psychonomic Bulletin & Review*, 5, 694–7.

Ambrosini, E., Costantini, M., and Sinigaglia, C. (2011) Grasping with the eyes. *Journal Of Neurophysiology*, 106(3), 1437–42.

Andrews, K. (2005) Chimpanzee theory of mind: Looking in all the wrong places? *Mind & Language*, 20, 521–36.

Andrews, K. (2012) *Do Apes Read Minds?* Cambridge, MA: MIT Press.

Anscombe, G. E. M. (1957) *Intention*. Oxford: Basil Blackwell.

Apperly, I. A. and Butterfill, S. A. (2009) Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4), 953–70.

Archibald, S. J., Mateer, C. A., and Kerns, K. A. (2001) Utilization behavior: Clinical manifestations and neurological mechanisms. *Neuropsychology Review*, 11, 117–30.

Armstrong, D. M. (2004) In defense of the cognitivist theory of perception. *The Harvard Review of Philosophy*, 12, 19–26.

Austin, J. L. (1964) *Sense and Sensibilia*. Oxford: Oxford University Press.

Bach, K. (1978) A representational theory of action. *Philosophical Studies*, 34, 361–79.

Ballard, D. H. (1996) On the function of visual representation. In Akins K. (ed.), *Perception*. New York: Oxford University Press, pp. 111–31.

Baron-Cohen, S. (1995) *Mindblindness*. Cambridge, MA: MIT Press.

Baron-Cohen, S. (2007) The Biology of the Imagination. In Wells, R. H. and McFadden, J. (eds), *Human Nature?* London: The Continuum International Publishing Group.

Bateson, M. (2002) Context-dependent foraging choices in risk-sensitive starlings. *Animal Behavior*, 64, 251–60.

Battaglia-Mayer, A. and Caminiti, R. (2002) Optic ataxia as a result of the breakdown of the global tuning fields of parietal neurones. *Brain*, 125, 225–37.

Batty, C. (2010) Scents and sensibilia. *American Philosophical Quarterly*, 47, 103–18.

Batty, C. (2011) Smelling lessons. *Philosophical Studies*, 153, 161–74.

Bayne, T. (2009) Perception and the reach of phenomenal content. *Philosophical Quarterly*, 59, 385–404.

Bechara, A., Damasio, H., Tranel, D., and Damasio, A. R. (1997) Deciding advantageously before knowing the advantageous strategy. *Science*, 275 (5304), 1293.

Bechara, A., Damasio, H., Tranel, D., and Damasio, A. R. (2005) The Iowa Gambling Task and the somatic marker hypothesis: Some questions and answers. *Trends in Cognitive Sciences*, 9(4), 159–62.

Bell, D. E., Raiffa, H., and Tversky, A. (1988) Descriptive, normative, and prescriptive interactions in decision making. In Bell, D. E., Raiffa, H., and Tversky, A. (eds), *Decision Making: Descriptive, Normative, and Prescriptive Interactions*. New York: Cambridge University Press, pp. 9–30.

Bermúdez, José Luis (2009) Mindreading in the animal kingdom. In Lurz R. (ed.), *The Philosophy of Animal Minds.* Cambridge: Cambridge University Press, pp. 89–107.

Berryhill, M. E., Phuong, L., Picasso, L., Cabeza, R., and Olson, I. R. (2007) Parietal lobe and episodic memory: Bilateral damage causes impaired free recall of autobiographical memory. *Journal of Neuroscience*, 27, 14415–23.

Bertelson, P. (1999) Ventriloquism: A case of cross-modal perceptual grouping. In Aschersleben, G., Bachmann, T., and Musseler, J. (eds), *Cognitive Contributions to the Perception of Spatial and Temporal Events*, Amsterdam: Elsevier, pp. 347–62.

Bertelson, P. and de Gelder, B. (2004) The psychology of multimodal perception. In Spence, C. and Driver, J. (eds), *Crossmodal Space and Crossmodal Attention*, Oxford: Oxford University Press, pp. 141–77.

Blaser, E., Pylyshyn, Z. W., and Holcombe, A. O. (2000) Tracking an object through feature space. *Nature*, 408, 196–9.

Block, N. (1995) A confusion about consciousness. *Behavioral and Brain Sciences*, 18, 227–47.

Bloom, P. (2010) *How Pleasure Works: The New Science of Why We Like What We Like*. New York: W. W. Norton & Co.

Bonini, L., Rozzi, S., Serventi, F. U., Simone, L., Ferrari, P. F., and Fogassi, L. (2010) Ventral premotor and inferior parietal cortices make distinct contribution to action organization and intention understanding. *Cerebral Cortex*, 20(6), 1372–85.

Botterill, G. (1996) Folk psychology and theoretical status. In Carruthers, P. and Smith, P. K. (eds), *Theories of Theories of Mind*. Cambridge: Cambridge University Press, pp. 105–18.

Brand, M. (1968) Danto on basic actions. *Nous*, 2, 187–90.

Brand, M. (1979) The fundamental question of action theory. *Nous*, 13, 131–51.

Brand, M. (1984) *Intending and Acting*. Cambridge, MA: MIT Press.

Bräuer, J., Call, J., and Tomasello, M. (2007) Chimpanzees really know what others can see in a competitive situation. *Animal Cognition*, 10, 439–48.

Bregman, A. S. (1990) *Auditory scene analysis: The perceptual organization of sound*. Cambridge, MA: Bradford Books, MIT.

Brewer, B. (2006) Perception and content. *European Journal of Philosophy*, 14, 165–81.

Brewer, B. (2007) Perception and its objects. *Philosophical Studies*, 132, 87–97.

Brewer, B. (2011) *Perception and its objects*. Oxford: Oxford University Press.

Bridgeman, B., Peery, S., and Anand, S. (1997) Interaction of cognitive and sensorimotor maps of visual space. *Perception & Psychophysics*, 59, 456–9.

Briscoe, R. (2008) Another look at the two visual systems hypothesis. *Journal of Conscious Studies*, 15, 35–62.

Briscoe, R. (2009) Egocentric spatial representation in action and perception. *Philosophy and Phenomenological Research*, 79, 423–60.

Briscoe, R. (2011) Mental imagery and the varieties of amodal perception. *Pacific Philosophical Quarterly*, 92, 153–73.

Brogaard, B. (2011) Are there unconscious perceptual processes? *Consciousness and Cognition*, 20, 449–63.

Brogaard, B. (forthcoming a) *Does Perception Have Content?* New York: Oxford University Press.

Brogaard, B. (forthcoming b) Conscious vision for action versus unconscious vision for action? *Cognitive Science*, 35, 1076–104.

Brooks, R. A. (1991) Intelligence without representation. *Artificial Intelligence*, 47, 139–59.

Bruno, N. (2001) When does action resist visual illusions? *Trends in Cognitive Sciences*, 5, 385–8.

Bruno, N. and Bernardis, P. (2002) Dissociating perception and action in Kanizsa's compression illusion. *Psychonomic Bulletin & Review*, 9, 723–30.

Bulloch, M. J., Boysen, S. T., and Furlong E. E. (2008) Visual attention and its relation to knowledge states in chimpanzees, Pan troglodytes. *Animal Behaviour*, 76, 1147–55.

Burge, T. (2005) Disjunctivism and perceptual psychology. *Philosophical Topics*, 33, 1–78.

Burge, T. (2010) *The Origins of Objectivity*. Oxford: Clarendon.

Butterfill, S. and Apperly, I. A. (forthcoming) How to construct a minimal theory of mind. *Mind & Language*.

Butterfill, S. and Sinigaglia, C. (forthcoming) Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*.

Byrne, A. (2007) Possibility and imagination. *Philosophical Perspectives*, 21, 125–44.

Byrne, A. and Logue, H. (2008) Either/or. In Haddock, A. and MacPherson, F. (eds), *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.

Byrne, P., Becker, S., Burgess, N. (2007) Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychological Review*, 114, 340–75.

Byrne, R. (1995) *The Thinking Ape*. Oxford: Oxford University Press.

Byrne, R. (1998) The early evolution of creative thinking: Evidence from monkeys and apes. In Mithen S. (ed.), *Creativity in Human Evolution and Prehistory*. London: Routledge, pp. 110–24.

Byrne, R. and Whiten, A. (1988) *Machiavellian Intelligence*. Oxford: Clarendon Press.

Call, J. and Tomasello, M. (2005) What chimpanzees know about seeing revisited: An explanation of the third kind. In Hoerl, C., Eilan, N., McCormack, T., and Roessler, J. (eds), *Joint Attention.* Oxford: Oxford University Press.

Call, J. and Tomasello, M. (2008) Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, 12, 187–92.

Campbell, J. (1993) The role of physical objects in spatial thinking. In Eilan, N., McCarthy R., and Brewer, B. (eds), *Spatial Representation*. Oxford: Blackwell, pp. 65–95.

Campbell, J. (1994) *Past, Space, and Self*. Cambridge, MA: MIT Press.

Campbell, J. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.

Campbell, J. (2011) Tyler Burge: Origins of objectivity. *Journal of Philosophy*, 108, 269–85.

Campbell, K. (1990) *Abstract Particulars*. Oxford: Blackwell.

Caraco, T. (1981) Energy budgets, risk, and foraging preferences in dark-eyed juncos (Junco hyemalis). *Behavioral Ecology and Sociobiology*, 8, 213–17.

Cardellicchio, P., Sinigaglia, C., and Costantini, M. (2013) Grasping affordances with the other's hand: A TMS study. *Social Cognitive and Affective Neuroscience*, 8, 455–9.

Carruthers, P. and Smith, P. K. (1996) *Theories of Theories of Mind.* Cambridge: Cambridge University Press.

Casati, R. and Dokic, J. (1994) *La Philosopie du Son*, Paris: Chambon.

Cassam, Q. (2007) *The Possibility of Knowledge*. Oxford: Oxford University Press.

Cavina-Pratesi, C., Gustav Kuhn, C., Ietswaart, M., and Milner, A. D. (2011) The magic grasp: Motor expertise in deception. *PLoS ONE*, 6(2), e16568.

Chalmers, D. J. (2002) Does conceivability entail possibility? In Gendler, T. S. and Hawthorne, J. (eds), *Conceivability and Possibility*. Oxford: Oxford University Press, pp. 145–200.

Chalmers, D. J. (2006) Perception and the fall from Eden. In Gendler, T. and Hawthorne, J. (eds), *Perceptual Experience*. New York: Oxford University Press, pp. 49–125.

Chemero, A. (2001) What we perceive when we perceive affordances. *Ecological Psychology*, 13, 111–16.

Chemero, A. (2009) *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Chen, M. K., Lakshminaryanan, V., and Santos, L. (2006) How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*, 114, 517–32.

Cheney, D. L. and Seyfarth, R. M. (1990) *How Monkeys See the World.* Chicago, IL: University of Chicago Press.

Cheney, D. L., Seyfarth, R. M., and Silk, J. B. (1995) The response of female baboons (Papio cynocephalus ursinus) to anomalus social interactions: Evidence for causal reasoning? *Journal of Comparative Psychology*, 109, 134–41.

Clark, A. (1995) Moving minds: Re-thinking representation in the heat of situated action. In Tomberlin, J. (ed.), *Philosophical Perspectives 9: AI, Connectionism and Philosophical Psychology*. Atascadero, CA: Ridgeview.

Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Clark, A. (2001) Visual experience and motor action: Are the bonds too tight? *Philosophical Review*, 110, 495–519.

Clark, A. (2009) Perception, action, and experience: Unraveling the golden braid. *Neuropsychologia*, 47, 1460–8.

Clark, A. (2000) *A Theory of Sentience*. Oxford: Clarendon Press.

Clark, A. (2004) Sensing, objects, and awareness: Reply to commentators. *Philosophical Psychology*, 17, 563–89.

Clark, A. (2011) Vicissitudes of non-visual objects. *Philosophical Studies*, 153, 175–81.

Clark, J. (2010) Relations of homology between higher order emotions and basic emotions. *Biology & Philosophy*, 25, 75–94.

Clarke, T. (1965) Seeing surfaces and physical objects. In Black, M. (ed.), *Philosophy in America*. Ithaca, NY: Cornell University Press, pp. 98–114.

Clayton, N. S., Griffiths, D. P., Emery, N. J., and Dickinson, A. (2001) Elements of episodic-like memory in animals. *Philosophical Transactions of the Royal Society B*, 356, 1483–91.

Cohen, J. (2004) Objects, places, and perception. *Philosophical Psychology*, 17, 471–95.

Costantini, M., Ambrosini, E., Tieri, G., Sinigaglia, C., and Committeri, G. (2010) Where does an object trigger an action? An investigation about affordances in space. *Experimental Brain Research*, 207, 95–103.

Costantini, M., Committeri, G., and Sinigaglia, C. (2011) Ready both to your and to my hands: Mapping the action space of others. *PLoS ONE*, 6(4), e17923.

Crane, T. (2006) Is there a perceptual relation? In Gendler, T. and Hawthorne, J. (eds), *Perceptual Experience*. Oxford: Oxford University Press, pp. 126–46.

Crane, T. (2009) Is perception a propositional attitude? *Philosophical Quarterly*, 59, 452–69.

Csibra, G. (2007) Action mirroring and action interpretation: An alternative account. In Haggard, P., Rosetti, Y., and Kawato, M. (eds), *Sensorimotor Foundations of Higher Cognition. Attention and Performance XXII*. Oxford: Oxford University Press, pp. 435–59.

Csibra, G. (2008) Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, 107, 705–17.

Csibra, G., Bíró, S., Koós, O., and Gergely, G. (2003) One-year-old infants use teleological representations of actions productively. *Cognitive Science*, 27, 111–33.

Csibra, G., Gergely, G., Bíró, S., Koós, O., and Brockbank, M. (1999) Goal attribution without agency cues: The perception of "pure reason" in infancy. *Cognition*, 72, 237–67.

Currie, G. (1995a) Visual imagery as the simulation of vision. *Mind and Language*, 10, 25–44.

Currie, G. (1995b) *Image and Mind: Film, Philosophy, and Cognitive Science*. Cambridge: Cambridge University Press.

Currie, G. and Ravenscroft, I. (1996) Mental simulation and motor imagery. *Philosophy of Science*, 64, 161–80.

Currie, G. and Ravenscroft, I. (2002) *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford: Oxford University Press.

Cusack, R., Decks, J., Aikman, G., and Carlyon, R. P. (2004) Effects of location, frequency region, and time course of selective attention on auditory scene analysis. *Journal of Experimental Psychology: Human Perception and Performance*, 30, 643–56.

Cussins, A. (2003) Experience, thought and activity. In Gunther, Y. (ed.), *Essays on Nonconceptual Content*. Cambridge, MA: MIT Press, pp. 133–63.

Cuthbert, B. N. et al. (2000) Brain potentials in affective picture processing: Covariation with autonomic arousal and affective report. *Biological Psychology*, 52, 95–111.

Damasio, A. R. (1994) *Descartes' Error*. New York: Putnam.

Danto, A. (1963) What we can do. *Journal of Philosophy*, 90, 435–45.

Danto, A. (1965) Basic actions. *American Philosophical Quarterly*, 2, 141–8.

Daprati, E. and Gentilucci, M. (1997) Grasping an illusion. *Neuropsychologia*, 35, 1577–82.

Darwall, S. (1998) Empathy, sympathy, care. *Philosophical Studies*, 89, 261–82.

Darwin, C. (1899) *The Expression of Emotion in Man and Animals*. New York: Appleton.

Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Oxford University Press.

Davidson, D. (1982) Rational animals. *Dialectica*, 36, 317–27.

Davidson, D. (2001) *Subjective, Intersubjective, Objective*. New York: Oxford University Press.

Davies, M. (1994) The mental simulation debate. In Peacocke, C. (ed.), *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind.* Oxford: Oxford University Press, pp. 99–127.

de Bruin, L., Strijbos, D., and Slors, M. (2011) Early social cognition: Alternatives to implicit mindreading. *Review of Philosophy and Psychology*, 2, 375–95.

de Sousa, R. (1987) *The Rationality of Emotions*. Cambridge, MA: MIT Press.

de Waal, F. B. M. (1982) *Chimpanzee Politics*. London: Jonathan Cape.

de Waal, F. B. M. (2009) *The Age of Empathy*. New York: McClelland.

Debes, R. (2010) Which empathy? Limitations in the mirrored "understanding" of emotion. *Synthese*, 175, 219–39.

Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., van de Moortele, P. F., and Le Bihan, D. (1998) Imaging unconscious semantic priming. *Nature*, 395, 597–600.

Della Sala, S., Marchetti, C., Spinnler, H. (1991) Right-sided anarchic (alien) hand: A longitudinal study. *Neuropsychologia*, 29, 1113–27.

Dennett, D. C. (1996) Seeing is believing—or is it? In Akins, K. (ed.), *Perception*. Oxford: Oxford University Press, pp. 111–31.

Deonna, J. (2006) Emotion, perception and perspective. *Dialectica*, 60, 29–46.

Deonna, J. (2007) The structure of empathy. *Journal of Moral Philosophy*, 4, 99–116.

Dijkerman, H. C. and De Haan, E. H. F. (2007) Somatosensory processes subserving perception and action. *Behavioral and Brain Sciences*, 30, 189–201.

Doggett, T. and Egan, A. (2007) Wanting things you don't want: The case for an imaginative analogue of desires. *Philosopher's Imprint*, 9, 1–17.

Dokic, J. (1998) The ontology of perception: Bipolarity and content. *Erkenntnis*, 48, 153–69.

Dolk, T., Hommel, B., Colzato, L. S., Schütz-Bosbach, S., Prinz, W., and Liepelt, R. (2011) How "social" is the Social Simon Effect? *Frontiers in Psychological Sciences*, 2, 84.

Dorrance, B. R. and Zentall, T. R. (2001) Imitative learning in Japanese quail (Coturnix japonica) depends on the motivational state of the observer quail at the time of observation. *Journal of Comparative Psychology*, 115(1), 62–7.

Dretske, F. (1969) *Seeing and Knowing*. London: Routledge.

Dretske, F. (1973) Perception and other minds. *Nous*, 7, 34–44.

Dretske, F. (2007) What change blindness teaches about consciousness. *Philosophical Perspectives*, 21, 215–30.

Dretske, F. (2010) On what we see. In Nanay, B. (ed.), *Perceiving the World*. New York: Oxford University Press, pp. 54–67.

Dreyfus, H. (2005) Overcoming the myth of the mental: How philosophers can profit from the phenomenology of everyday expertise. *Proceedings and Addresses of the American Philosophical Association*, 79(2), 47–65.

Ellis, R. and Tucker, M. (2000) Micro-affordance: The potentiation of components of action by seen objects. *British Journal of Psychology*, 91, 451–71.

Emery, N. J. and Clayton, N. S. (2004) The mentality of crows: Convergent evolution of intelligence in corvids and apes. *Science*, 306, 1903–7.

Evans, G. (1982) *The Varieties of Reference*. Oxford: Oxford University Press.

Evans, J. St. B. T. (2007) *Hypothetical Thinking: Dual Processes in Reasoning and Judgement*. Hove: Psychology Press.

Evans, J. St. B. T. and Over, D. (2004) *If*. Oxford: Oxford University Press.

Feagin, S. L. (1996) *Reading with Feeling: The Aesthetics of Appreciation*. Ithaca, NY: Cornell University Press.

Feingold, G. A. (1915) The influence of suggestion on imagination. *American Journal of Psychology*, 26, 540–9.

Findlay, J. M. and Gilchrist, I. D. (2003) *Active Vision: The Psychology of Looking and Seeing*. Oxford: Oxford University Press.

Flombaum, J. I. and Santos, L. R. (2005) Rhesus monkeys attribute perceptions to others. *Current Biology*, 15, 447–52.

Fodor, J. A. (1981) *Representations*. Cambridge, MA: MIT Press.

Fodor, J. A. (1983) *The Modularity of Mind*. Cambridge, MA: MIT Press.

Fodor, J. A. and Pylyshyn, Z. (1981) How direct is visual perception? Some reflections on Gibson's "ecological approach". *Cognition*, 9, 139–96.

Fox, C. R. and See, K. E. (2003) Belief and preference in decision under uncertainty. In Hardman, D. and Macchi, L. (eds), *Thinking: Psychological Perspectives on Reasoning, Judgment, and Decision Making*. New York: Wiley, pp. 273–314.

Fox, C. R. and Tversky, A. (1998) A belief-based account of decision under uncertainty. *Management Science*, 44, 879–95.

Franz, V. (2001) Action does not resist visual illusions. *Trends in Cognitive Sciences*, 5, 457–9.

Franz, V. (2003) Manual size estimation: A neuropsychological measure of perception? *Experimental Brain Research*, 151, 471–7.

Franz, V., Bülthoff, H. H., and Fahle, M. (2003) Grasp effects of the Ebbinghaus illusion: Obstacle avoidance is not the explanation. *Experiential Brain Research*, 149, 470–7.

Franz, V. and Gegenfurtner, K. (2008) Grasping visual illusions: Consistent data and no dissociation. *Cognitive Neuropsychology*, 25, 920–50.

Franz, V., Gegenfurtner, K., Bülthoff, H., and Fahle, M. (2000) Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science*, 11, 20–5.

Frith, C. D., Blakemore, S. J., and Wolpert D. M. (2000) Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London, B Biological Sciences*, 355, 1771–88.

Funkhouser, E. (2006) The determinable–determinate relation. *Nous*, 40, 548–69.

Gallagher, S. (2001) The practice of mind: Theory, simulation, or interaction? *Journal of Consciousness Studies*, 8, 83–107.

Gallagher, S. (2005) *How the Body Shapes the Mind*. Oxford: Oxford University Press.

Gallagher, S. (2007a) Logical and phenomenological arguments against simulation theory. In Hutto, D. and Ratcliffe, M. (eds), *Folk Psychology Re-assessed*. Dordrecht: Springer, pp. 63–78.

Gallagher, S. (2007b) Simulation trouble. *Social Neuroscience*, 2, 353–65.

Gallagher, S. (2008) Direct perception in the intersubjective context. *Consciousness and Cognition*, 17, 535–43.

Gallagher, S. and Meltzoff, A. N. (1996) The earliest sense of self and others. *Philosophical Psychology*, 9, 213–36.

Gallese, V. (2005) "Being like me": Self–other identity, mirror neurons and empathy. In Hurley, S. and Chater, N. (eds), *Perspectives on Imitation I*. Cambridge, MA: MIT Press, pp. 101–18.

Gallese, V. (2007) Before and below "theory of mind": Embodied simulation and the neural correlates of social cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 659–69.

Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996) Action recognition in the premotor cortex. *Brain*, 119, 593–609.

Gallese, V. and Goldman, A. (1998) Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Science*, 3, 493–501.

Gallese, V., Keysers, C., and Rizzolatti, G. (2004) A unifying view of the basis of social cognition. *Trends in Cognitive Sciences*, 8, 396–403.

Gaut, B. (1998) Imagination, interpretation, and film. *Philosophical Studies*, 89, 331–41.

Gaut, B. (1999) Identification and emotion in narrative film. In Plantinga, C. and Smith, G. M. (eds), *Passionate Views: Thinking about Film and Emotion*. Baltimore, MD: Johns Hopkins University Press, pp. 200–16.

Gawronski, B. and Bodenhausen, G. V. (2006) Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132(5), 692–731.

Gawronski, B., Strack, F., and Bodenhausen, G. V. (2009) Attitudes and cognitive consistency: The role of associative and propositional processes. In Petty, R. E., Fazio, R. H., and Brinol, P. (eds), *Attitudes: Insights from the New Implicit Measures*. Mahwah, NJ: Erlbaum, pp. 85–117.

Gendler, T. S. (2008a) Alief in action (and reaction). *Mind & Language*, 23, 552–85.

Gendler, T. S. (2008b) Alief and belief. *Journal of Philosophy*, 105(10), 634–63.

Gendler, T. S. (2011) *Intuition, Imagination, and Philosophical Methodology*. Oxford: Oxford University Press.

Gentilucci, M., Cheiffe, S., Daprati, E., Saetti, M. C., and Toni, I. (1996) Visual illusion and action. *Neuropsychologia*, 34, 369–76.

Gentilucci, M., Daprati, E., Toni, I., Chieffi, S., and Saetti, M. C. (1995) Unconscious updating of grasp motor program. *Experimental Brain Research*, 105, 291–303.

Gergely, G., Nadasdy, Z., Csibra, G., and Biro, S. (1995) Taking the intentional stance at 12 months of age. *Cognition*, 56, 165–93.

Gibson, J. J. (1966) *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton Mifflin.

Gibson, J. J. (1979) *An Ecological Approach to Visual Perception*, Boston, MA: Houghton Mifflin.

Gillam, B. (1998) Illusions at century's end. In Hochberg, J. (ed.), *Perception and Cognition at Century's End*. San Diego, CA: Academic Press, pp. 95–136.

Giovannetti, T., Buxbaum, L. J., Biran, I., and Chatterjee, A. (2005) Reduced endogenous control in alien hand syndrome: Evidence from naturalistic action. *Neuropsychologia*, 43, 75–88.

Gobbini, M. I., Koralek, A. C., Bryan, R. E., Montgomery K. J., and Haxby J. V. (2007) Two takes on the social brain: A comparison of theory of mind tasks. *Journal of Cognitive Neuroscience*, 19, 1803–14.

Goldie, P. (1999) Understanding other people's emotions. *Mind and Language*, 14, 394–423.

Goldie, P. (2000) *The Emotions: A Philosophical Exploration*. Oxford: Oxford University Press.

Goldie, P. (2002) Emotion, personality and simulation. In Goldie, P. (ed.), *Understanding Emotions: Mind and Morals*. Aldershot: Ashgate Publishing, pp. 97–109.

Goldie, P. (2004) *On Personality*. London: Routledge.

Goldie, P. (2007) Seeing what is the kind thing to do: Perception and emotion in morality. *Dialectica*, 35, 347–61.

Goldman, A. (1992) In defense of the simulation theory. *Mind and Language*, 7, 104–19.

Goldman, A. (2000) The mentalizing folk. In Sperber, D. (ed.), *Metarepresentations*. Oxford: Oxford University Press.

Goldman, A. (2006) *Simulating Minds: The Philosophy, Psychology and Neuroscience of Mindreading*. New York: Oxford University Press.

Gómez, J-C. (1996) Non-human primate theories of (non-human primate) minds: Some issues concerning the origins of mind-reading. In Carruthers, P., Smith, P. K. (eds), *Theories of Theories of Mind*. Cambridge: Cambridge University Press, pp. 330–43.

Gómez, J-C. (2004) *Apes, monkeys, children, and the growth of mind*. Cambridge, MA: Harvard University Press.

Gómez, J-C. (2008) The evolution of pretence: From intentional availability to intentional non-existence. *Mind & Language*, 23, 586–606.

Gonzalez, C., Ganel, T., Whitwell, R., Morrissey, B., and Goodale, M. (2008) Practice makes perfect, but only with the right hand: Sensitivity to perceptual illusions with awkward grasps decreases with practice in the right but not the left hand. *Neuropsychologia*, 46, 624–31.

Goodale, M. A. (2011) Transforming vision into action. *Vision Research*, 51, 1567–87.

Goodale, M. A. and Milner, A. D. (2004) *Sights Unseen*. Oxford: Oxford University Press.

Goodale, M. A., Pelisson, D., and Prablanc, C. (1986) Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320, 748–50.

Goodale, M. A. and Humphrey, G. K. (1998) The objects of action and perception. *Cognition*, 67, 181–207.

Goodall, J. (1971) *In the Shadow of Man*. Boston, MA: Houghton Mifflin.

Gordon, R. M. (1995a) Sympathy, simulation, and the impartial spectator. *Ethics*, 105. 727–42.

Gordon, R. M. (1995b) Simulation without introspection or inference from me to you. In Davies, M. and Stone, T. (eds), *Mental Simulation*. Oxford: Blackwell, pp. 53–67.

Greenspan, P. (1988) *Emotions and Reason*. New York: Routledge.

Griffiths, P. (1997) *What Emotions Really Are*. Chicago, IL: University of Chicago Press.

Griffiths, P. (2004) Is "emotion" a natural kind? In Solomon, R. C. (ed.), *Philosophers on Emotion*. New York: Oxford University Press, pp. 233–49.

Grush, R. (2004) The emulation theory of representation: Motor control, imagery and perception. *Behavioral and Brain Sciences*, 27, 377–442.

Guagnanoa, D., Ruscoli, E., and Umilta, C. A. (2010) Sharing a task or sharing space? On the effect of the confederate in action coding in a detection task. *Cognition*, 114, 348–55.

Haffenden, A. and Goodale, M. A. (1998) The effect of pictorial illusion on prehension and perception. *Journal of Cognitive Neuroscience*, 10, 122–36.

Haffenden, A., Schiff, K. C., and Goodale, M. A. (2001) The dissociation between perception and action in the Ebbinghaus illusion: Nonillusory effects of pictorial cues on grasp. *Current Biology*, 11, 177–81.

Haggard, P. (2005a) Conscious intention and motor cognition. *Trends in Cognitive Sciences*, 9, 290–5.

Haggard, P. (2005b) Conscious intention and the sense of agency. In Sebanz, N. (ed.), *Disorders of Volition*. Oxford: Oxford University Press.

Haggard, P. and Clark, S. (2003) Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition*, 12, 695–707.

Hamlin, J. K., Wynn, K., and Bloom, P. (2007) Social evaluation by preverbal infants. *Nature*, 450, 557–60.

Hare, B., Call, J., Agnetta, B., and Tomasello, M. (2000) Chimpanzees know what conspecifics do and do not see. *Animal Behavior*, 59, 771–85.

Hare, B., Call, J., and Tomasello, M. (2001) Do chimpanzees know what conspecifics know? *Animal Behavior*, 61, 139–51.

Hare, B., Call, J., and Tomasello, M. (2006) Chimpanzees deceive a human by hiding. *Cognition*, 101, 495–514.

Hatzimoysis, A. (2007) The case against unconscious emotions. *Analysis*, 67, 292–9.

Heal, J. B. (1995) How to think about thinking. In Davies, M. and Stone, T. (eds), *Mental Simulation*. Oxford: Blackwell, pp. 33–52.

Heal, J. B. (1998) Co-cognition and off-line simulation: Two ways of understanding the simulation approach. *Mind & Language*, 13, 477–98.

Heath, C. and Tversky, A. (1991) Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4, 5–28.

Heidegger, M. (1977) The origin of the work of art. In Krell, D. F. (ed.), *Basic Writings: From Being and Time (1927) to The Task of Thinking (1964)*. San Francisco, CA: Harper Collins.

Held, R. (1965) Plasticity in the sensory-motor system. *Scientific American*, 213, 84–94.

Hellie, B. (forthcoming) *Conscious Life*. Oxford: Oxford University Press.

Hess, W. R., Bürgi, S., and Bucher, V. (1946) Motorische funktion des tectal und tegmentalgebietes. *Monatschrift für psychiatrische Neurologie*, 72, 571–9.

Heyes, C. M. (1994) Reflection on self-recognition in primates. *Behavioral and Brain Sciences*, 16, 524–5.

Heyes, C. M. (1998) Theory of mind in nonhuman primates. *Behavioral and Brain Sciences*, 21, 101–34.

Heyes, C. M. (2001) Causes and consequences of imitation. *Trends in Cognitive Sciences*, 5, 253–61.

Heyes, C. M. (2005) Imitation by association. In Hurley, S. and Chater, N. (eds), *Perspectives on Imitation*. Cambridge, MA: MIT Press.

Hickock, G. and Poeppel, D. (2007) The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8, 393–401.

Hommel, B., Colzato, L. S., and Van den Wildenberg, W. P. M. (2009) How social are task representations? *Psychological Science*, 20, 794–8.

Hommel, B., Müsseler, J., Aschersleben, G., and Prinz, W. (2001) The theory of event coding: A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849–931.

Hopkins, R. (2012) What Perky did not show. *Analysis*, 72, 431–9.

Hornsby, J. (2004) Agency and alienation. In de Caro, M. and Macarthur, D. (eds), *Naturalism in Question*. Cambridge, MA: Harvard University Press.

Horridge, G. A., Zhang, S-W., O'Carrol, D. (1992) Insect perception of illusory contours. *Philosophical Transactions of the Royal Society London, B*, 337, 59–64.

Hsee, C. K. and Weber, E. U. (1999) Cross-national differences in risk preference and lay predictions. *Journal of Behavioral Decision Making*, 12(2), 165–79.

Humberstone, I. L. (1992) Direction of fit. *Mind*, 101, 59–83.

Hume, D. (1739/1978) *A Treatise of Human Nature*. Oxford: Oxford University Press.

Humphreys, G. W. and Bedford, J. (2011) The relations between joint action and theory of mind: A neuropsychological analysis. *Experimental Brain Research*, 211, 357–69.

Humphreys, G. W. and Riddoch, M. J. (2001) Detection by action: Neuropsychological evidence for action-defined templates in search. *Nature Neuroscience*, 4, 84–8.

Hurley, S. L. (1998) *Consciousness in Action*. Cambridge, MA: Harvard University Press.

Hurly, T. and Ossen, M. (1999) Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behavior*, 58, 59–66.

Hustert, R. and Baldus, M. (2010) Ballistic movements of jumping legs implemented as variable components of cricket behavior. *Journal of Experimental Biology*, 213, 4055–64.

Hutto, D. D. (2004) The limits of spectatorial folk psychology. *Mind & Language*, 19, 548–73.

Hutto, D. D. (2007) Folk psychology without theory or simulation. In Hutto, D. and Ratcliffe, M. (eds), *Folk Psychology Re-assessed,* Dordrecht: Springer, pp. 115–35.

Hutto, D. D. (2008) *Folk Psychological Narratives: The Sociocultural Basis of Understanding Reasons*. Cambridge, MA: MIT Press.

Hutto, D. D. (2011) Elementary mind minding, enactivist-style. In Seemann, A. (ed.), *Joint Attention*. Cambridge, MA: MIT Press, pp. 307–41.

Hutto, D. D. and Myin, E. (2013) *Radicalizing Enactivism*. Cambridge, MA: MIT Press.

Hutto, D. D. and Ratcliffe, M. (2007) *Folk Psychology Re-assessed*. Dordrecht: Springer.

Ingmanson, E. J. (2002) Empathy in a bonobo. In Mitchell, R. W. (ed.), *Pretending and Imagination in Animals and Children*. Cambridge: Cambridge University Press, pp. 280–4.

Ishida, H., Nakajima, K., Inase, M., and Murata, A. (2009) Shared mapping of own and others' bodies in visuotactile bimodal area of monkey parietal cortex. *Journal of Cognitive Neuroscience*, 22, 83–96.

Ishiguro, H. (1967) Imagination. *Proceedings of the Aristotelian Society, Supplementary Volume*, 41, 37–56.

Israel, D., Perry, J., and Tutiya, S. (1993) Executions, motivations and accomplishments. *Philosophical Review*, 102, 515–40.

Jackson, F. (1977) *Perception: A representative theory*. Cambridge: Cambridge University Press.

Jackson, S. and Shaw, A. (2000) The Ponzo illusion affects grip-force but not grip-aperture scaling during prehension movements. *Journal of Experimental Psychology HPP*, 26, 418–23.

Jacob, P. (2005) Grasping and perceiving an object. In Brooks, A. and Akins, K. (eds), *Cognition and the Brain*. Cambridge: Cambridge University Press, pp. 241–86.

Jacob, P. (2008) What do mirror neurons contribute to human social cognition? *Mind and Language*, 23, 190–223.

Jacob, P. (2009) The tuning-fork model of human social cognition: A critique. *Consciousness and Cognition*, 18, 229–43.

Jacob, P. and Jeannerod, M. (2003) *Ways of Seeing: The Scope and Limits of Visual Cognition*. Oxford: Oxford University Press.

Jagnow, R. (2011) Ambiguous figures and spatial contents of perceptual experience. *Philosophy and the Cognitive Sciences*, 10, 325–46.

James, W. (1884) What is an emotion? *Mind*, 9, 188–205.

James, W. (1890) *The Principles of Psychology*. New York: Holt.

Jeannerod, M. (1988) *The Neural and Behavioral Organization of Goal-Directed Movements*. New York: Oxford University Press.

Jeannerod, M. (1994) The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, 17, 187–245.

Jeannerod, M. (1997) *The Cognitive Neuroscience of Action*. Oxford: Blackwell.

Jeannerod, M. and Jacob, P. (2005) Visual cognition: A new look at the two-visual systems model. *Neuropsychologia*, 43, 301–12.

Jeannerod, M. and Pacherie, E. (2004) Agency, simulation, and self-identification. *Mind and Language*, 19(2), 113–46.

Johnston, W. E. (1921) *Logic: Part I*. Cambridge: Cambridge University Press.

Johnson, M. H., Dziurawiec, S., Ellis, H., and Morton, J. (1991) Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40, 1–19.

Kaas, J. H. and Hackett, T. A. (1999) "What" and "where" processing in auditory cortex. *Nature Neuroscience*, 2, 1045–7.

Kahneman, D. and Tversky, A. (1973) Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207–33.

Kahneman, D. and Tversky, A. (1979) Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–91.

Kakei, S., Hoffman, D. S., and Strick, P. L. (2001) Direction of action is represented in the ventral premotor cortex. *Nature Neuroscience*, 4, 1020–5.

Kamewari, K., Kato, M., Kanda, T., Ishiguro, H., and Hiraki, K. (2005) Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion. *Cognitive Development*, 20, 303–20.

Kaminski, J., Call, J., and Tomasello, M. (2004) Body orientation and face orientation: Two factors controlling apes' begging behavior from humans. *Animal Cognition*, 7, 216–23.

Kaminski, J., Call, J., and Tomasello, M. (2008) Chimpanzees know what others know, but not what they believe. *Cognition*, 109, 224–34.

Kaplan, D. (1977/1989) Demonstratives. In Almog, J., Perry, J., and Wettstein, H. (eds), *Themes from Kaplan*. Oxford: Oxford University Press, pp. 481–564.

Karin-D'Arcy, M. and Povinelli, D. J. (2002) Do chimpanzees know what each other see? A closer look. *International Journal of Comparative Psychology*, 15, 21–54.

Kawato, M. (1999) Internal models for motor control and trajectory planning. *Current Opinions in Neurobiology*, 9, 718–27.

Keil, A. et al. (2002) Large-scale neural correlates of affective picture viewing. *Psychophysiology*, 39, 641–9.

Kelly, S. (2010) The normative nature of perceptual experience. In Nanay, B. (ed.), *Perceiving the World*. New York: Oxford University Press, pp. 146–59.

Keysers, C. and Gazzola, V. (2007) Integrating simulation and theory of mind: From self to social cognition. *Trends in Cognitive Sciences*, 11, 194–6.

Kind, A. (2001) Putting the image back to imagination. *Philosophy and Phenomenological Research*, 62, 85–109.

Koffka, K. (1935) *Principles of Gestalt Psychology*. New York: Harcourt Brace.

Kosslyn, S. M. (1980) *Image and Mind*. Cambridge, MA: Harvard University Press.

Kosslyn, S. M. (1994) *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, MA: MIT Press.

Kosslyn, S. M., Thompson, W. L., and Ganis, G. (2006) *The Case for Mental Imagery*. Oxford: Oxford University Press.

Kosslyn, S. M., Thompson, W. L., Kim, I. J., and Alpert, N. M. (1995) Topographical representations of mental images in primary visual cortex. *Nature*, 378, 496–8.

Kravitz, D. J., Saleem, K. I., Baker, C. I., and Mishkin, M. (2011) A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*, 12, 217–30.

Kriegel, U. (2007) The phenomenologically manifest. *Phenomenology and the Cognitive Sciences*, 6, 115–36.

Króliczak, G., Heard, P., Goodale, M. A., and Gregory, R. L. (2006) Dissociation of perception and action unmasked by the hollow-face illusion. *Brain Research*, 1080, 9–16.

Kubovy, M. and Valkenburg, D. V. (2001) Auditory and visual objects. *Cognition*, 80, 97–126.

Kuhbandner, C., Pekrun, R., and Maier, M. A. (2010) The role of positive and negative affect in the "mirroring" of other persons' actions. *Cognition and Emotion*, 24, 1182–90.

Kuhlmeier, V., Wynn, K., and Bloom, P. (2003) Attribution of dispositional states by 12-month-olds. *Psychological Science*, 14, 402–8.

Kulvicki, J. (2008) The nature of noise. *Philosophers' Imprint*, 8, 1–16.

Lakatos, I. (1970) Falsification and the methodology of scientific research programmes. In Lakatos, I. and Musgrave, A. (eds), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.

Lakatos, I. (1974) Science and pseudoscience. In Vesey, G. (ed.), *Philosophy in the Open*. Milton Keynes: Open University Press, pp. 96–102.

Lakshminarayanan, V., Chen, M. K., and Santos, L. R. (2011) The evolution of decision-making under risk: Framing effects in monkey risk preferences. *Journal of Experimental Social Psychology*, 47, 689–93.

Lamme, V. A. F. (2003) Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7, 12–18.

Lau, L-Y. and Ranyard, R. (2005) Chinese and English probabilistic thinking and risk taking in gambling. *Journal of Cross-Cultural Psychology*, 36(5), 621–7.

Leddington, J. (2013) What we hear. In Brown, R. (ed.), *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*. Springer Studies in Brain and Mind. Dordrecht: Springer, in press.

Lee, D. (2006) Neural basis of quasi-rational decision making. *Current Opinions in Neurobiology*, 16, 191–8.

Leslie, A. M. (1987) Pretense and representation in infancy: The origins of the theory of mind. *Psychological Review*, 94(4), 412–26.

Lewis, C. I. (1929) *Mind and the World-Order: Outline of a Theory of Knowledge.* New York: Charles Scribner's Sons.

Lewis, J. W. and Van Essen, D. C. (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *Journal of Comparative Neurology*, 428, 112–37.

Lhermitte, F. (1983) Utilization behavior and its relation to lesions of the frontal lobes. *Brain*, 106, 237–55.

Libet, B. (1985) Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8, 529–66.

Lopes, L. L. (1987) Between hope and fear: The psychology of risk. In Berkowitz, L. (ed.), *Advances in Experimental Social Psychology, Volume 20*. New York: Academic Press, pp. 255–95.

Lopes, D. M. (2003) *The aesthetics of photographic transparency*. Mind, 112, 433–48.

Luo, Y. (2011) Three-month-old infants attribute goals to a non-human agent. *Developmental Science*, 14, 453–60.

Lycan, W. (1996) *Consciousness and Experience*. Cambridge, MA: MIT Press.

Lycan, W. (2000) The slighting of smell. In Bhushan, N. and Rosenfeld, S. (eds), *Of Minds and Molecules: New Philosophical Perspectives on Chemistry*. Oxford: Oxford University Press, pp. 273–89.

Mack, A. and Rock, I. (1998) *Inattentional Blindness*. Cambridge, MA: MIT Press.

Macpherson, F. (2011) Cross-modal experiences. *Proceedings of the Aristotelian Society*, 111, 429–68.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., Thorn, T. M. J., and Castelli, L. (1997) On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*, 33, 471–89.

Mandik, P. (2005) Action oriented representation. In Brook, A. and Akins, K. (eds), *Cognition and the Brain: The Philosophy and Neuroscience of Movement*. Cambridge: Cambridge University Press, pp. 284–305.

Marsh, B. and Kacelnik, A. (2002) Framing effects and risky decisions in starlings. *Proceedings from the National Academy of Sciences*, 99, 3352–5.

Martin, M. G. F. (1997) The shallows of the mind. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 71, 75–98.

Martin, M. G. F. (2002a) The transparency of experience. *Mind & Language*, 17, 376–425.

Martin, M. G. F. (2002b) Particular thoughts and singular thought. *Royal Institute of Philosophy Supplement*, 51, 173–214.

Martin, M. G. F. (2004) The limits of self-awareness. *Philosophical Studies*, 120, 37–89.

Martin, M. G. F. (2006) On being alienated. In Gendler, T. and Hawthorne, J. (eds), *Perceptual Experience*. Oxford: Oxford University Press, pp. 354–410.

Martin, M. G. F. (forthcoming) *Uncovering Appearances*. Oxford: Oxford University Press.

Masrour, F. (2011) Is perceptual phenomenology thin? *Philosophy and Phenomenological Research*, 83, 366–97.

Matthen, M. (2004) Features, places and things. *Philosophical Psychology*, 17, 497–519.

Matthen, M. (2005) *Seeing, Doing, and Knowing: A Philosophical Theory of Sense-Perception*. Oxford: Clarendon Press.

Matthen, M. (2010a) How things look (and what things look that way). In Nanay, B. (ed.), *Perceiving the World*. New York: Oxford University Press, pp. 226–53.

Matthen, M. (2010b) On the diversity of auditory objects. *Review of Philosophy and Psychology*, 1, 63–89.

Matthen, M. (forthcoming) Visual demonstratives. In Raftopoulos, A. and Machamer, P. (eds), *Perception, Realism, and the Problem of Reference*. Cambridge: Cambridge University Press.

Maynard, P. (2012) What's so funny? Comic content in depiction. In Meskin, A. and Cook, R. T. (eds), *The Art of Comics: A Philosophical Approach*. London: Wiley-Blackwell.

McCoy, A. and Platt, M. (2005) Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, 8(9), 1220.

McIntosh, R. D. and Lashleya, G. (2008) Matching boxes: Familiar size influences action programming. *Neuropsychologia*, 46, 2441–4.

McKay, R. T. and Dennett, D. (2009) The evolution of misbelief. *Behavioral and Brain Sciences*, 32, 493–510.

McNeill, W. E. S. (2012) On seeing that someone is angry. *European Journal of Philosophy*, 20, 575–97.

Melis, A. P., Call, J., and Tomasello, M. (2006) Chimpanzees conceal visual and auditory information from others. *Journal of Comparative Psychology*, 120, 154–62.

Mellers, B. A., Schwartz, A., Ho, K., and Ritov, I. (1997) Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science*, 8, 423–9.

Mellet, E., Petit, L., Mazoyer, B., Denis, M., and Tzourio, N. (1998) Reopening the mental imagery debate: Lessons from functional anatomy. *NeuroImage*, 8, 129–39.

Meltzoff, A. N. (1995) Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31, 838–50.

Millikan, R. G. (1995) Pushmi-pullyu representations. *Philosophical Perspectives*, IX, 185–200.

Millikan, R. G. (2004) *Varieties of Meaning*. Cambridge, MA: MIT Press.

Milner, A. D. and Goodale, M. A. (1995) *The Visual Brain in Action*. Oxford: Oxford University Press.

Milner, A. D. and Goodale, M. A. (2008) Two visual systems re-viewed. *Neuropsychologia*, 46, 774–85.

Mitchell, R. W. (2002) *Pretending and Imagination in Animals and Children*. Cambridge: Cambridge University Press.

Mithen, S. (2001) The evolution of imagination: An archeological perspective. *SubStance*, 30, pp. 28–54.

Moll, H. and Tomasello, M. (2006) Level 1 perspective-taking at 24 months of age. *British Journal of Developmental Psychology*, 24, 603–13.

Morgado N., Muller D., Gentaz E., and Palluel-Germain R. (2011) Close to me? The influence of affective closeness on space perception. *Perception*, 40(7), 877–9.

Morsanyi, K. and Handley, S. J. (2012) Logic feels so good-I like it! Evidence for intuitive detection of logicality in syllogistic reasoning. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 38, 596–616.

Mulligan, K. (1999) Perception, particulars and predicates. In Fisette, D. (ed.), *Consciousness and Intentionality: Models and Modalities of Attribution*. Dordrecht: Kluwer, pp. 163–94.

Mulligan, K., Smith, B., and Simons, P. (1984) Truth makers. *Philosophy and Phenomenological Research*, 44, 278–321.

Nanay, B. (2004) Taking twofoldness seriously: Walton on imagination and depiction. *Journal of Aesthetics and Art Criticism*, 62, 285–9.

Nanay, B. (2005) Is twofoldness necessary for representational seeing? *British Journal of Aesthetics*, 45, 263–72.

Nanay, B. (2006a) Does what we want influence what we see? In *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006)*. Hillsdale, NJ: Lawrence Erlbaum, pp. 615–21.

Nanay, B. (2006b) Perception, action and identification in the theatre. In Krasner, D. and Staltz, D. (eds), *Staging Philosophy*. Ann Arbor, MI: Michigan University Press, pp. 244–54.

Nanay, B. (2007) Four theories of amodal perception. In McNamara, D. S. and Trafton, J. G. (eds), *Proceedings of the 29th Annual Conference of the Cognitive Science Society (CogSci 2007)*. Hillsdale, NJ: Lawrence Erlbaum, pp. 1331–6.

Nanay, B. (2008) Picture perception and the two visual subsystems. In Love, B. C., McRae, K., and Sloutsky, V. M. (eds), *Proceedings of the 30th Annual Conference of the Cognitive Science Society (CogSci 2008)*. Hillsdale, NJ: Lawrence Erlbaum, pp. 975–80.

Nanay, B. (2009a) Imagining, recognizing and discriminating: Reconsidering the ability hypothesis. *Philosophy and Phenomenological Research*, 79, 699–717.

Nanay, B. (2009b) The properties of singular causation. *The Monist*, 72, 112–32.

Nanay, B. (2010a) Attention and perceptual content. *Analysis*, 70, 263–70.

Nanay, B. (2010b) A modal theory of function. *Journal of Philosophy*, 107, 412–31.

Nanay, B. (2010c) Perception and imagination: Amodal perception as mental imagery. *Philosophical Studies*, 150, 239–54.

Nanay, B. (2010d) Philosophy of perception: The new wave. In Nanay, B. (ed.), *Perceiving the World*. New York: Oxford University Press, pp. 3–12.

Nanay, B. (2010e) Adam Smith's concept of sympathy and its contemporary interpretations. *Adam Smith Review*, 5, 85–105.

Nanay, B. (2010f) Inflected and uninflected perception of pictures. In Abell, C. and Bantilaki, K. (eds), *Philosophical Perspectives on Depiction*. Oxford: Oxford University Press, pp. 181–207.

Nanay, B. (2010g) Rational reconstruction reconsidered. *The Monist*, 93, 595–615.

Nanay, B. (2010h) Population thinking as trope nominalism. *Synthese*, 177, 91–109.

Nanay, B. (2010i) Three ways of resisting racism. *The Monist*, 93, 256–82.

Nanay, B. (2011a) Do we see apples as edible? *Pacific Philosophical Quarterly*, 92, 305–22.

Nanay, B. (2011b) Do we sense modalities with our sense modalities? *Ratio*, 24, 299–310.

Nanay, B. (2011c) Perceiving pictures. *Phenomenology and the Cognitive Sciences*, 10, 461–80.

Nanay, B. (2011d) Popper's Darwinian analogy. *Perspectives on Science*, 19, 337–54.

Nanay, B. (2011e) Ambiguous pictures, attention and perceptual content. *Phenomenology and the Cognitive Sciences*, 10, 557–61.

Nanay, B. (2011f) Three ways of resisting essentialism about natural kinds. In Campbell, J. K. and Slater, M. H. (eds), *Carving Nature at its Joints: Topics in Contemporary Philosophy*, *Volume 8*. Cambridge, MA: MIT Press, pp. 175–97.

Nanay, B. (2011g) What if reality has no architecture? *The Monist*, 94, 181–97.

Nanay, B. (2012a) Action-oriented perception. *European Journal of Philosophy*, 20, 430–46.

Nanay, B. (2012b) The philosophical implications of the Perky experiments. *Analysis*, 72, 439–43.

Nanay, B. (2012c) Singularist semirealism. *British Journal for the Philosophy of Science*, 64, 371–94.

Nanay, B. (2012d) Perceiving tropes. *Erkenntnis*, 77, 1–14.

Nanay, B. (2012e) Perceptual phenomenology. *Philosophical Perspectives*, 26, 235–46.

Nanay, B. (2012f) Function attribution depends on the explanatory context. *Journal of Philosophy*, 109, 623–7.

Nanay, B. (forthcoming a) *Aesthetics as Philosophy of Perception*. Oxford: Oxford University Press.

Nanay, B. (2013) Success semantics: The sequel. *Philosophical Studies* 165: 151–165.

Nanay, B. (forthcoming b) Philosophy versus literature: Against the Discontinuity Thesis. *Journal of Aesthetics and Art Criticism*.

Nanay, B. (forthcoming c) Empirical problems with anti-representationalism. In Brogaard, B. (ed.), *Does Perception have Content?* New York: Oxford University Press.

Nanay, B. (forthcoming d) Naturalizing action theory. In Sprevak, M. and Kallestrup, J. (eds), *New Waves in the Philosophy of Mind*. Hampshire: Palgrave Macmillan.

Nanay, B. (forthcoming f) The representationalism versus relationalism debate: Explanatory contextualism about perception. *European Journal of Philosophy*.

Nanay, B. (forthcoming g) Naturalism and experimental philosophy. In Fischer, E. (ed.), *Philosophical Insights*. Oxford: Oxford University Press.

Nanay, B. (ms) Perceptual content and the content of mental imagery.

Nehaniv, C. L. and Dautenhahn, K. (2002) The correspondence problem. In Dautenhahn, K. and Nehaniv, C. L. (eds), *Imitation in Animals and Artifacts*, Cambridge, MA: MIT Press, 41–61.

Neill, A. (1996) Empathy and (film) fiction. In Bordwell, D. and Carroll, N (eds), *Post Theory: Reconstructing Film Studies*. Madison, WI: Wisconsin University Press, pp. 175–94.

Neiworth, J. J. (1992) Cognitive aspects of movement estimation: A test of imagery in animals. In Honig, W. K. and Fetterman, J. G. (eds), *Cognitive Aspects of Stimulus Control*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 323–46.

Nettle, D. (2001) *Strong Imagination*. Oxford: Oxford University Press.

Newell, B. R and Shanks, D. R. (forthcoming) Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*.

Nichols, S. and Stich, S. (2003) *Mindreading*. Oxford: Oxford University Press.

Nieder, A. (2002) Seeing more than meets the eye: Processing of illusory contours in animals. *Journal of Comparative Physiology A*, 188, 294–60.

Nieder, A. and Wagner, H. (1999) Perception and neuronal coding of subjective contours in the owl. *Nature Neuroscience*, 2, 660–3.

Noë, A. (2002) Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9(5–6), 1–12.

Noë, A. (2003) Causation and perception: The puzzle unravelled. *Analysis*, 63, 93–100.

Noë, A. (2004) *Action in Perception*. Cambridge, MA: MIT Press.

Noordhof, P. (2002) Imagining objects and imagining experiences. *Mind & Language*, 17, 426–55.

Norman, J. (2002) Two visual systems and two theories of perception. *Behavioral and Brain Sciences*, 25, 73–144.

Nudds, M. (2001) Experiencing the production of sounds. *European Journal of Philosophy*, 9, 210–29.

Nudds, M. (2010) What are auditory objects? *Review of Philosophy and Psychology*, 1, 105–22.

Nudds, M. and O'Callaghan, C. (2009) *Sounds and Perception*. Oxford: Oxford University Press.

O'Callaghan, C. (2007) *Sounds*. Oxford: Oxford University Press.

O'Callaghan, C. (2008a) Perception and multimodality. In Margolis, E., Samuels, R., and Stich, S. (eds), *Oxford Handbook to Philosophy and Cognitive Science*. Oxford: Oxford University Press.

O'Callaghan, C. (2008b) Seeing what you hear: Crossmodal illusions and perception. *Philosophical Issues*, 18, 316–38.

O'Callaghan, C. (2009) Sounds. In Bayne, T., Cleermans, A., and Wilken, P. (eds), *The Oxford Companion to Consciousness*. Oxford: Oxford University Press.

O'Callaghan, C. (2011) Lessons from beyond vision (sounds and audition). *Philosophical Studies*, 153, 143–60.

O'Connell, S. M. (1995) Empathy in chimpanzees: Evidence for theory of mind? *Primates*, 36, 397–410.

O'Neill, D. K. (1996) Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development*, 67, 659–77.

O'Regan, K. (1992) Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46, 461–88.

O'Shaugnessy, B. (1980) *The Will*. Cambridge: Cambridge University Press.

Oakley, D. A. (1985) Cognition and imagery in animals. In Oakley, D. A. (ed.), *Brain and Mind*. London: Methuen, pp. 99–131.

Oatley, K. and Johnson-Laird, P. N. (1987) Towards a cognitive theory of emotions. *Cognition and Emotion*, 1, 29–50.

Olson, M. A. and Fazio, R. H. (2006) Reducing automatically activated racial prejudice through implicit evaluative conditioning. *Personality and Social Psychology Bulletin*, 32(4), 421.

Onishi, K. H. and Baillargeon, R. (2005) Do 15-month-old infants understand false beliefs? *Science*, 308, 255–8.

Ortony, A. and Turner, W. (1990) What's basic about basic emotions? *Psychological Review*, 97, 315–31.

Pacherie, E. (2000) The content of intentions. *Mind and Language*, 15, 400–32.

Pacherie, E. (2007) The anarchic hand syndrome and utilization behavior: A window onto agentive self-awareness. *Functional Neurology*, 22, 211–17.

Pacherie, E. (2011) Nonconceptual representations for action and the limits of intentional control. *Social Psychology*, 42, 67–73.

Page, J. W., Duhamel, P., and Crognale, M. A. (2011) ERP evidence of visualization at early stages of visual processing. *Brain and Cognition*, 75(2), 141–6.

Papineau, D. and Heyes, C. M. (2006) Rational or associative? Imitation in Japanese quail. In Nudds, M. and Hurley, S. (eds), *Rational Animals*. Oxford: Oxford University Press.

Parr, L. A. (2001) Cognitive and physiological markers of emotional awareness in chimpanzees, *Pan troglodytes*. *Animal Cognition*, 4, 223–9.

Parr, L. A. and Hopkins, W. D. (2000) Brain temperature asymmetries and emotional perception in chimpanzees, *Pan troglodytes*. *Physiology and Behavior*, 71, 363–71.

Parr, L. A., Winslow, J. T., and Davis, M. (2002) Somatic and cardiac responses to startle in differentially reared rhesus monkeys (*Macaca mulatta*). *Behavioral Neuroscience*, 116, 378–86.

Pasnau, R. (1999) What is sound? *Philosophical Quarterly*, 49, 309–24.

Patten, W. (1920) The grand strategy of evolution. Boston, MA: R. G. Badger.

Patterson, F. G. P. and Cohn, R. H. (1994) Self-recognition and self-awareness in lowland gorillas. In Parker, G. (ed.), *Self-awareness in Animals and Humans*. Cambridge: Cambridge University Press, pp. 273–290.

Paulignan, Y., MacKenzie, C. L., Marteniuk, R. G., and Jeannerod, M. (1991) Selective perturbation of visual input during prehension movements: 1. The effect of changing object position. *Experimental Brain Research*, 83, 502–12.

Pautz, A. (2010) An argument for the intentional view of visual experience. In Nanay, B. (ed.), *Perceiving the World.* Oxford: Oxford University Press, pp. 254–309.

Pavani, F., Boscagli, I., Benvenuti, F., Rabuffetti, M., and Farnè, A. (1999) Are perception and action affected differently by the Titchener circles illusion? *Experimental Brain Research*, 127, 95–101.

Peacocke, C. (1983) *Sense and Content*. New York: Oxford University Press.

Peacocke, C. (1985) Imagination, experience, and possibility: A Berkeley-an view defended. In Foster, J. and Robinson, H. (ed.), *Essays on Berkeley*. Oxford: Clarendon, pp. 19–35.

Peacocke, C. (1986) Analogue content. *Proceedings of the Aristotelian Society Supplementary Volume*, 60, 1–17.

Peacocke, C. (1989) Perceptual content. In Almong, J., Perry, J., and Wettstein, H. (eds), *Themes from Kaplan*. Oxford: Oxford University Press, pp. 297–329.

Peacocke, C. (1992) *A Study of Concepts*. Cambridge, MA: MIT Press.

Peacocke, C. (2007) Mental action and self-awareness (I). In McLaughlin, B. and Cohen, J. (eds), *Contemporary Debates in Philosophy of Mind*. Oxford: Blackwell, pp. 358–76.

Pelisson, D., Prablanc, C., Goodale, M. A., and Jeannerod, M. (1986) Visual control of reaching movements without vision of the limb: II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. *Experimental Brain Research*, 62, 303–11.

Penn, D. C. and Povinelli, D. J. (2007) On the lack of evidence that non-human animals possess anything remotely resembling a "theory of mind". *Philosophical Transactions of the Royal Society, B*, 362, 731–44.

Pereboom, D. (2004) Is our conception of agent causation coherent? *Philosophical Topics*, 32, 275–86.

Perkins, M. (1983) *Sensing the World*. Indianapolis, IN: Hackett Publishing.

Perky, C. W. (1910) An experimental study of imagination. *American Journal of Psychology*, 21, 422–52.

Perner, J. (1991) *Understanding the Representational Mind*. Cambridge, MA and London: MIT Press and Bradford Books.

Perner, J. (1996) Simulation as explication of prediction-implicit knowledge about the mind: Arguments for a simulation-theory mix. In Carruthers, P. and Smith, P. K. (eds), *Theories of Theories of Mind*. Cambridge: Cambridge University Press, pp. 90–104.

Perner, J. and Ruffman, T. (2005) Infants' insight into the mind: How deep? *Science*, 308, 214–16.

Perry, J. (2001) *Knowledge, Possibility and Consciousness*. Cambridge, MA: MIT Press.

Petrinovich, L. and O'Neill, P. (1996) Influence of wording and framing effects on moral intuitions. *Ethology and Sociobiology*, 17, 145–71.

Pitcher, G. (1965) Emotions. *Mind*, 74, 324–46.

Platt, M. and Glimcher, P. (1999) Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–8.

Platts, M. (1979) *Ways of Meaning*. London: Routledge and Kegan Paul.

Podzebenko, K., Egan, G. F., and Watson, J. D. G. (2002) Widespread dorsal stream activation during a parametric mental rotation task, revealed with functional magnetic resonance imaging. *NeuroImage*, 15, 547–58.

Poincaré, H. (1905/1958) *The Value of Science*. New York: Dover.

Port, R. and van Gelder, T. J. (1995) *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge MA: MIT Press.

Porter, J., Craven, B., Khan, R., Chang, S-J., Kang, I., Judkewitz, B., and Sobel, N. (2007) Mechanisms of scent tracking in humans. *Nature Neuroscience*, 10, 27–9.

Posner, M. I. (1980) Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 2–25.

Posner, M. I. (1984) Current research in the study of selective attention. In Donchin, E. (ed.), *Cognitive Psychophysiology: Event Related Potentials and the Study of Cognition.* Hillsdale, NJ: Erlbaum.

Posner, M. I., Walker, J. A., Friedrich, F. J., and Rafal, R. D. (1984) Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*, 4, 1863–74.

Povinelli, D. J. and Eddy, T. J. (1996) What young chimpanzees know about seeing. *Monographs in Social Research in Child Development*, 61, 1–152.

Povinelli, D. J. and Vonk, J. (2003) Chimpanzee minds: Suspiciously human? *Trends in Cognitive Science*, 7, 157–60.

Povinelli, D. J. and Vonk, J. (2004) We don't need a microscope to explore the chimpanzee mind. *Mind and Language*, 19, 1–28.

Premack, D. and Woodruff, G. (1978) Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1, 515–26.

Premack, D. and Premack, A. J. (1997) Infants attribute value G to the goal-directed actions of self-propelled objects. *Journal of Cognitive Neuroscience*, 9, 848–56.

Preston, S. and de Waal, F. (2002) Empathy: Its ultimate and proximate bases. *Behavioral and Brain Sciences*, 25, 1–20.

Prinz, J. J. (2004) *Gut Reactions*. New York: Oxford University Press.

Prinz, J. J. (2006) Beyond appearances. In Gendler, T. and Hawthorne, J. (eds), *Perceptual Experience*. Oxford: Oxford University Press, pp. 434–60.

Prinz, J. J. (2010) How do perceptual states become conscious? In Nanay, B. (ed.), *Perceiving the World: New Essays on Perception*. New York: Oxford University Press, pp. 310–32.

Prosser, S. (2011) Affordances and the phenomenal character in spatial perception. *Philosophical Review*, 120, 475–513.

Pulvermuller, F. and Hauk, O. (2005) Functional links between motor and language systems. *European Journal of Neuroscience*, 21, 793–7.

Pylyshyn, Z. W. (1999) Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–65.

Pylyshyn, Z. W. (2002) Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, 25, 157–82.

Pylyshyn, Z. W. (2003) Return of the mental image: Are there really pictures in the brain? *Trends in Cognitive Sciences*, 7, 113–18.

Pylyshyn, Z. W. (2007) *Things and Places: How the Mind Connects with the World*. Cambridge, MA: MIT Press.

Quine, W. V. O. (1984) Reply to Putnam. In Hahn, L. E. and Schillp, P. A. (eds), *The Philosophy of W. V. Quine*. La Salle: Open Court.

Quoidbach, J., Gilbert, D. T., and Wilson, T. D. (2013) The end of history illusion. *Science*, 339, 96–8.

Ramachandran, V. S. (1987) Visual perception of surfaces: A biological theory. In Petry, S. and Meyer, G. (eds), *The Perception of Illusory Contours*. Berlin: Springer, pp. 93–108.

Ramsey, W. M. (2007) *Representation Reconsidered*. Cambridge: Cambridge University Press.

Raftopoulos, A. (forthcoming) Late vision: Processes and epistemic status. *Frontiers in Perception Science*.

Ratcliffe, M. J. (2007) *Rethinking Commonsense Psychology: A Critique of Folk Psychology, Theory of Mind and Simulation*. Basingstoke: Palgrave Macmillan.

Ravenscroft, I. (1998) What is it like to be someone else? Simulation and empathy. *Ratio*, 11, 170–85.

Reaux, J., Theall, L., and Povinelli, D. J. (1999) A longitudinal investigation of chimpanzees' understanding of visual perception. *Child Development*, 70, 275–90.

Recanati, F. (ms) 'Imagining de se', Unpublished.

Reed, C. L., Klatzky, R. L., and Halgren, E. (2005) What vs. where in touch: An fMRI study. *NeuroImage*, 25, 718–26.

Regolin, L. and Vallortigara, G. (1995) Perception of partly occluded objects by young chicks. *Perception and Psychophysics*, 57, 971–6.

Reynolds, S. L. (1989) Imagining oneself to be another. *Nous*, 23, 615–33.

Rice, G. E. J. (1964) Aiding behavior vs. fear in the albino rat. *Psychological Record*, 14, 165–70.

Richardson, A. (1969) *Mental Imagery*. New York: Springer.

Riddoch, M. J., Edwards, M. G., Humphreys, G. W., West, R., and Heafield, T. (1998) Visual affordances direct action: Neuropsychological evidence from manual interference. *Cognitive Neuropsychology*, 15, 645–93.

Rilling, M. E. and Neiworth, J. J. (1987) Theoretical and methodological considerations for the study of imagery in animals. *Learning and Motivation*, 18, 57–79.

Rizzolatti, G. and Matelli, M. (2003) Two different streams form the dorsal visual system: Anatomy and functions. *Experimental Brain Research*, 153, 146–57.

Rizzolatti, G. and Sinigaglia, C. (2008) *Mirrors in the Brain: How our Minds Share Actions and Emotions.* New York: Oxford University Press.

Rottenstreich, Y. and Hsee, C. K. (2001) Money, kisses, and electric shocks: On the affective psychology of risk. *Psychological Science*, 12, 185–90.

Rowlands, M. (2006) *Body Language*: *Representation in Action*. Cambridge, MA: MIT Press.

Rozin, P., Millman, L., and Nemeroff, C. (1986) Operation of the laws of sympathetic magic in disgust and other domains. *Journal of Personality and Social Psychology*, 50, 703–12.

Rozzi, S., Ferrari, P. F., Bonini, L., Rizzolatti, G., and Fogassi, L. (2008) Functional organization of inferior parietal lobule convexity in the macaque monkey: Electrophysiological characterization of motor, sensory and mirror responses and their correlation with cytoarchitectonic areas. *European Journal of Neuroscience*, 28, 1569–88.

Ryle, G. (1949) *The Concept of Mind*. London: Hutchinson.

Schellenberg, S. (2010) Particularity and phenomenology of perceptual experience. *Philosophical Studies*, 149, 19–48.

Schellenberg, S. (forthcoming) Belief and desire in imaginations and immersions. *Journal of Philosophy*.

Schenk, T. and McIntosh, R. D. (2010) Do we have independent visual streams for perception and action? *Cognitive Neuroscience*, 1, 52–78.

Schnall, S., Benton, J., and Harvey, S. (2008) With a clean conscience: Cleanliness reduces the severity of moral judgments. *Psychological Science*, 19, 1219–22.

Schnall, S., Zadra, J. R. and Proffitt, D. R. (2010) Direct evidence for the economy of action: Glucose and the perception of geographic slant. *Perception*, 39, 464–82.

Scholl, B. J. and Tremoulet, P. D. (2000) Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4, 299–309.

Schupp, H. T. et al. (2000) Affective picture processing: The late positive potential is modulated by motivational relevance. *Psychophysiology*, 37, 257–61.

Schupp, H. T. et al. (2004) Brain processes in emotional perception: Motivated attention. *Cognition and Emotion*, 18, 593–611.

Schwitzgebel, E. (2008) The unreliability of naive introspection. *Philosophical Review*, 117, 245–73.

Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Sebanz, N., Bekkering, H., and Knoblich, G. (2006) Joint actions: Bodies and minds moving together. *Trends in Cognitive Sciences*, 10, 70–6.

Sebanz, N., Knoblich, G., and Prinz, W. (2003) Representing others' actions: Just like one's own? *Cognition*, 88(3), B11–21.

Sebanz, N., Knoblich, G., and Prinz, W. (2005a) How two share a task: Corepresenting stimulus–response mappings. *Journal of Experimental Psychology: Human Perception and Performance*, 31(6), 1234–46.

Sebanz, N., Knoblich, G., Stumpf, L., and Prinz, W. (2005b) Far from action blind: Representation of others' actions in individuals with autism. *Cognitive Neuropsychology*, 22, 433–54.

Segal, S. J. (1972) Assimilation of a stimulus in the construction of an image: The Perky effect revisited. In Sheehan, P. W. (ed.), *The Function and Nature of Imagery*. New York: Academic Press, pp. 203–30.

Segal, S. J. and Nathan, S. (1964) The Perky effect: Incorporation of an external stimulus into an imagery experience under placebo and control conditions. *Perceptual and Motor Skills*, 19, 385–95.

Seidler, R. D., Noll, D. C., and Thiers, G. (2004) Feedforward and feedback processes in motor control. *NeuroImage*, 22, 1775–83.

Senju, A., Southgate, V., White, S., and Frith, U. (2009) Mindblind eyes: An absence of spontaneous theory of mind in Asperger syndrome. *Science*, 325, 883–5.

Shallice, T. and Burgess, P. W. (1991) Deficits in strategy application following frontal lobe damage in man. *Brain*, 114, 727–41.

Shallice, T., Burgess, P. W., Schon, F., and Baxter, D. M. (1989) The origins of utilization behaviour. *Brain*, 112, 1587–98.

Shams, L., Kamitani, Y., and Shimojo, S. (2000) What you see is what you hear. *Nature*, 408, 788.

Shepard, R. N. and Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science*, 171, 701–3.

Shoemaker, S. (1990) Qualities and qualia: What's in the mind. *Philosophy and Phenomenological Research*, 50, 109–31.

Siegal, M. and Varley, R. (2002) Neural systems involved in "theory of mind". *Nature Reviews Neuroscience*, 3, 463–71.

Siegel, S. (2002) Review of Austen Clark, *A Theory of Sentience*. *Philosophical Review*, 111, 135–8.

Siegel, S. (2005) The phenomenology of efficacy. *Philosophical Topics*, 33, 65–84.

Siegel, S. (2006a) Which properties are represented in perception? In Gendler, T. and Hawthorne, J. (eds), *Perceptual Experience*. Oxford: Oxford University Press, pp. 481–503.

Siegel, S. (2006b) Subject and object in the contents of visual experience. *Philosophical Review*, 115, 355–88.

Siegel, S. (2007) How can we discover the contents of experience? *Southern Journal of Philosophy* (Supp), 45, 127–42.

Siegel, S. (2009) The visual experience of causation. *Philosophical Quarterly*, 59, 519–40.

Siegel, S. (2010a) Do visual experiences have contents? In Nanay, B. (ed.), *Perceiving the World*. Oxford: Oxford University Press, pp. 333–68.

Siegel, S. (2010b) *The Contents of Visual Experience*. New York: Oxford University Press.

Siegel, S. (forthcoming) Affordances and the contents of perception. In Brogaard, B. (ed.), *Does Perception have Content?* New York: Oxford University Press.

Simmons, D. J. and Chabris, C. F. (1999) Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28, 1059–74.

Simon, J. R. (1990) The effects of an irrelevant directional cue on human information processing. In Proctor, R. W. and Reeve, T. G. (eds), *Stimulus–Response Compatibility: An Integrated Perspective. Advances in Psychology*, 65. Amsterdam: North-Holland, pp. 31–86.

Sinigaglia, C. (2009) Mirror in action. *Journal of Consciousness Studies*, 16, 309–34.

Slezak, P. (1987) Intending and acting (book review). *Journal of Philosophy*, 84, 49–54.

Slezak, P. (1989) How NOT to naturalize the theory of action. In Selzak, P. and Arbury, W. R. (eds), *Computers, Brains and Minds*. Dordrecht: Kluwer, pp. 137–66.

Sloman, S. A. (1996) The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.

Slovic, P., Fischhoff, B., and Lichtenstein, S. (1976) Cognitive processes and societal risk taking. In Carroll, J. S. and Payne, J. W. (eds), *Cognition and social behavior*. Hillsdale, NJ: Erlbaum, pp. 165–84.

Smith, A. (1759/1976) *The Theory of Moral Sentiments*. Raphael, D. D. and Macfie, A. L. (eds). Oxford: Clarendon Press.

Smith, J. (2006) Bodily awareness, imagination and the self. *European Journal of Philosophy*, 14, 49–68.

Smith, J. (2010) Seeing other people. *Philosophy and Phenomenological Research*, 81, 731–48.

Smith, M. A. (1987) The Humean Theory of Motivation. *Mind*, 96, 36–61.

Smith, M. (1997) Imagining from the inside. In Allen, R. and Smith, M. (eds), *Film Theory and Philosophy*. Oxford: Oxford University Press, pp. 412–30.

Song, H-J., Baillargeon, R., and Fisher, C. (2005) Can infants attribute to an agent a disposition to perform a particular action? *Cognition*, 98, 45–55.

Soteriou, M. (2000) The particularlity of visual perception. *European Journal of Philosophy*, 8, 173–89.

Southgate, V., Senju, A., and Csibra, G. (2007) Action anticipation through attribution of false belief by 2-year-olds. *Psychological Science*, 18, 587–92.

Spence, C. and Driver, J. (2004) *Crossmodal Space and Crossmodal Attention*. Oxford: Oxford University Press.

Spener, M. (2011) Using first person data about consciousness. *Journal of Consciousness Studies*, 18, 165–79.

Spener, M. and Bayne, T. (2010) Introspective humility. *Philosophical Issues*, 20, 1–22.

Stalnaker, R. (1976) Possible worlds. *Nous*, 10, 65–75.

Stazicker, J. (2011) Attention, visual consciousness and indeterminacy. *Mind & Language*, 26, 156–84.

Stein, B. E., Stanford, T. R., Wallace, M. T., Vaughan, J. W., and Jiang, W. (2004) Crossmodal spatial interactions in subcortical and cortical circuits. In Spence, C. and Driver, J. (eds), *Crossmodal Space and Crossmodal Attention*. Oxford: Oxford University Press, pp. 25–50.

Sterelny, K. (2003) *Thought in a Hostile World*. Oxford: Blackwell.

Steuber, K. R. (2006) *Rediscovering Empathy*. Cambridge, MA: MIT Press.

Stone, T. and Davies, M. (1996) The mental simulation debate: A progress report. In Carruthers, P. and Smith, K. P. (eds), *Theories of Theories of Mind*. Cambridge: Cambridge University Press, pp. 119–37.

Strawson, P. (1979) Perception and its objects. In MacDonald, G. F. (ed.), *Perception and Identity*: *Essays Presented to A. J. Ayer with His Replies*. Ithaca, NY: Cornell University Press, pp. 41–60.

Surian, L., Caldi, S., and Sperber, D. (2007) Attribution of beliefs by 13-month-old infants. *Psychological Science*, 18, 580–6.

Swain, S., Alexander, J., and Weinberg, J. M. (2008) The instability of philosophical intuitions. *Philosophy and Phenomenological Research*, 76, 138–55.

Tai, K., Zheng, X., and Narayanan, J. (2011) Touching a teddy bear mitigates negative effects of social exclusion to increase prosocial behavior. *Social Psychological and Personality Science*, 2, 618–26.

Thompson, B. (2009) Senses for senses. *Australasian Journal of Philosophy*, 87, 99–117.

Tom, S., Fox, C., Trepel, C., and Poldrack, R. (2007) The neural basis of loss aversion in decision making under risk. *Science*, 315, 515–18.

Tomasello, M. (1999) *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.

Tomasello, M. and Call, J. (1997) *Primate Cognition*. Oxford: Oxford University Press.

Tomasello, M. and Call, J. (2006) Do chimpanzees know what others see—or only what they are looking at? In Hurley, S. L. and Nudds, M. (eds), *Rational Animals?* Oxford: Oxford University Press.

Tomasello, M., Call, J., and Hare, B. (2003) Chimpanzees understand psychological states—The question is which ones and to what extent. *Trends in Cognitive Science*, 7, 153–6.

Tomasello, M., Carpenter, M., Call, J., Behne, T., and Moll, H. (2005) Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675–735.

Tooby, J. and Cosmides, L. (2005) Conceptual foundations of evolutionary psychology. In Buss, D. (ed.), *Handbook of Evolutionary Psychology*. Hoboken, NJ: John Wiley & Sons, pp. 5–67.

Travis, C. (2004) The silence of the senses. *Mind*, 113, 57–94.

Trehub, A. (1991) *The Cognitive Brain*. Cambridge, MA: MIT Press.

Trepel, C., Fox, C. R., and Poldrack, R. (2005) Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cognitive Brain Research*, 23, 34–50.

Tsai, C. C., Kuo, W. J., Hung, D. L., and Tzeng, O. J. (2008) Action corepresentation is tuned to other humans. *Journal of Cognitive Neuroscience*, 20(11), 2015–24.

Turvey, M. T. (1992) Affordances and prospective control: An outline of the ontology. *Ecological Psychology*, 4(3), 173–87.

Tversky, A. and Kahneman, D. (1981) The framing of decisions and the psychology of choice. *Science*, 211, 453–8.

Tversky, A. and Kahneman, D. (1982) Judgments of and by representativeness. In Kahneman, D., Slovic, P., and Tversky, A. (eds), *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.

Tversky, B. and Hard, B. M. (2009) Embodied and disembodied cognition: Spatial perspective-taking. *Cognition*, 110, 124–9.

Tye, M. (1991) *The Imagery Debate*. Cambridge, MA: MIT Press.

Tye, M. (2007) Intentionalism and the problem of no common content. *Philosophical Perspectives*, 21, 495–541.

Umiltà, M. A. et al. (2008) How pliers become fingers in the monkey motor system. *Proceedings of the National Academy of Sciences, USA*, 105, 2209–13.

Valdesolo, P. and DeSteno, D. (2006) Manipulations of emotional context shape moral judgment. *Psychological Science*, 17, 476–7.

Van Leeuwen, N. (2011) Imagination is where the action is. *Journal of Philosophy*, 108(2), 55–77.

Velleman, D. (1996) Self to self. *Philosophical Review*, 105, 39–76.

Velleman, D. (2000) *The Possibility of Practical Reason*. Oxford: Oxford University Press.

Vishton, P. (2004) Human vision focuses on information relevant to a task, to the detriment of information that is not relevant. *Behavioral and Brain Sciences*, 27, 53–4.

Vishton, P. and Fabre, E. (2003) Effects of the Ebbinghaus illusion on different behaviors. *Spatial Vision*, 16, 377–92.

Vroomen, J., Bertelson, P., and de Gelder, B. (2001) Auditory-visual spatial interactions: Automatic versus intentional components. In de Gelder, B., de Haan, E., and Heywood, C. (eds), *Out of Mind*. Oxford: Oxford University Press, pp. 140–50.

Wagenaar, W. (1988) *Paradoxes of Gambling Behaviour*. London: Erlbaum.

Wakker, P. P. (2004) On the composition of risk preference and belief. *Psychological Review*, 111, 236–41.

Walker, M. B. (1992) *The Psychology of Gambling*. Oxford: Pergamon.

Walton, K. (1990) *Mimesis and Make-Believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press.

Walton, K. (2002) Depiction, perception, and imagination. *The Journal of Aesthetics and Art Criticism*, 60, 27–35.

Wang, D. and Brown, G. J. (2006) *Computational Auditory Scene Analysis*. London: Wiley.

Watanabe, S. and Ono, K. (1986) An experimental analysis of "empathic" response: Effects of pain reactions of pigeon upon other pigeon's operant behavior. *Behavioural Processes*, 13, 269–77.

Watkins, S., Shams, L., Tanaka, S., Haynes, J. D., and Rees, G. (2006) Sound alters activity in human V1 in association with illusory visual perception. *Neuro-Image*, 31, 1247–56.

Wegner, D. (2002) *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

Wellman, H. M., Cross, D., and Watson, J. (2001) Meta-analysis of theory of mind development: The truth about false belief. *Child Development*, 72, 655–84.

Wheeler, M. (2005) *Reconstructing the Cognitive World*. Cambridge, MA: MIT Press.

Williams, B. (1973) Imagination and the self. In Williams, B., *Problems of the Self*, Cambridge: Cambridge University Press, pp. 26–45.

Williams, L. E. and Bargh, J. A. (2008) Experiencing physical warmth promotes interpersonal warmth. *Science*, 322, 606–7.

Wimmer, H. and Perner, J. (1983) Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13, 103–28.

Winkielman, P. and Berridge, K. C. (2004) Unconscious emotions. *Current Direction in Psychological Science*, 13, 120–3.

Winslow, J. T., Parr, L. A., and Davis, M. (2002) Acoustic startle, prepulse inhibition and fear-potentiated startle measured in rhesus monkeys. *Biological Psychiatry*, 51, 859–66.

Whiten, A. (1994) Grades of mindreading. In Lewis, C. and Mitchell, P. (eds), *Children's Early Understanding of Mind*. Hove: Erlbaum, pp. 47–70.

Whiten, A. (1996) When does smart behaviour-reading become mind-reading? In Carruthers, P. and Smith, P. K. (eds), *Theories of Theories of Mind*. Cambridge: Cambridge University Press, pp. 277–92.

Whiten, A. and Byrne R. W. (1997) *Machiavellian Intelligence II*. Cambridge: Cambridge University Press.

Witt, J. T., Proffitt, D. R., and Epstein, W. (2005) Tool use affects perceived distance, but only when you intend to use it. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 880–8.

Wollheim, R. (1973) Imagination and identification. In Wollheim, R., *On Art and the Mind: Essays and Lectures*. London: Allen Lane, pp. 54–83.

Wollheim, R. (1974) Identification and Imagination. In Wollheim, R. (ed.), *Freud: A Collection of Critical Essays*. New York: Anchor Press, pp. 172–95.

Wollheim, R. (1980) Seeing-as, seeing-in, and pictorial representation. In *Art and its Object*. Second Edition. Cambridge: Cambridge University Press, pp. 205–26.

Wollheim, R. (1987) *Painting as an Art*. Princeton, NJ: Princeton University Press.

Wollheim, R. (1998) On pictorial representation. *Journal of Aesthetics and Art Criticism*, 56, 217–26.

Yaari, M. E. (1987) The dual theory of choice under risk. *Econometrica*, 55, 95–117.

Yeshurun, Y. and Carrasco, M. (1998) Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, 396, 72–5.

Zahavi, D. (2008) Simulation, projection and empathy. *Consciousness and Cognition*, 17, 514–22.

Zajonc, R. B. (1984) On the primacy of affect. *American Psychologist*, 39, 117–23.

Zentall, T. R. (2003) Imitation by animals: How do they do it? *Current Directions in Psychological Science*, 12, 91–5.

Zhong, C-B. and Liljenquist, K. (2006) Washing away your sins: Threatened morality and physical cleansing. *Science*, 313, 1451–2.

Index