

Perceiving What Could Be. A Critique of Perception-Based Epistemologies of Possibility

MSc Thesis (*Afstudeerscriptie*)

written by

Alex Silviu Stan

(born 18th of November, 1999 in Cluj-Napoca, Romania)

under the supervision of **Dr. Tom Schoonen**, and submitted to the
Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense: **Members of the Thesis Committee:**

21st of February, 2025

Aybüke Özgün

Dr. Paul Dekker

Dr. Karolina Krzyzanowska



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis is a critique of the claim that perception can lead us to knowledge of possibilities and an attempt to offer an alternative view, based on the knowledge of our own abilities. First, I discuss and analyze available, perception-based epistemologies of possibility and I argue that all of them fail. I identify three types of issues with those accounts: they face *the problem of modal epistemic friction* and collapse into ability-based epistemologies of possibility, they run into circularity problems, or their proposed version of modal knowledge falls short of full-blown, actual world knowledge of possibility. Second, I propose a new ability-based epistemology of possibility, grounded in bodily awareness, that manages to avoid the problems that arise in the case of perception-based accounts and of the other available ability-based modal epistemology, namely the one proposed by Barbara Vetter. In particular, in contrast to Vetter's theory, the account that I propose can explain the phenomenon of *intrapersonal imitation*, which I argue is crucial in the process of ability knowledge acquisition.

Acknowledgements

I would first of all like to thank my supervisor, Tom Schoonen, for the guidance he offered me throughout this thesis building process. Even though our philosophical intuitions are very similar, my conversations with him pushed me to find my own thoughts and his critical feedback motivated me to clarify and develop them. His advice also made me a better writer, as I learned to keep the reader in mind more and guide them through the text more carefully.

I would also like to mention my friends and my MoL companions. Special thanks to Ștefan Zotescu, best friend, flat mate and fellow MoL student, for all the conversations and jokes we shared over the kitchen table and for the emotional support he offered me throughout our friendship. Those things kept me sane and motivated during the MoL coursework and during the thesis writing process. I would also like to thank the 'Bijlmer Gang', made up of Ștefan Zotescu, Sarah Dukic, Paulius Skaisgiris, Joel Saarinen and Lamia Tawam. Dinners and board game nights with them will always remain with me and they helped in navigating the Amsterdam social life, which can be difficult at times. I also want to thank my family, for always supporting me in my endeavours.

Finally, my biggest thanks to Toni, my partner. All the video calls, late night talks and dreams about the future kept me moving forward and brought happiness, hope and connection in difficult times. Without them none of this would have been possible.

Contents

1	Introduction	5
1.1	Modal Epistemology	5
1.2	The Problem of Modal Epistemic Friction	8
1.3	Thesis Aims	9
2	Climbing Trees and Reaching Mugs	12
2.1	Introduction	12
2.2	Strohminger’s Account	12
2.3	Critique	15
2.4	Conclusion	20
3	Pragmatic Representations	21
3.1	Introduction	21
3.2	Action and Pragmatic Representations	21
3.3	Do We Represent Objects as Q-able in Perception?	23
3.4	PMEF	31
3.5	Conclusion	32
4	Affordance Perception and Possibility	33
4.1	Introduction	33
4.2	Affordances	33
4.3	Recent Approaches	35
4.4	Circularity and PMEF	37
4.4.1	Two Options	38
4.4.2	A way out?	40
4.5	Conclusion	41
5	Ability Knowledge	42
5.1	Introduction	42
5.2	Abilities	43
5.3	An Agency-based Epistemology	45
5.3.1	Why Other Theories Fail	46
5.3.2	Vetter’s account	49
5.4	Bodily Awareness	53

5.5 Conclusion	56
6 Conclusion	57

Chapter 1

Introduction

1.1 Modal Epistemology

We are well accustomed to knowing truths about the world we live in. They are statements like 'The sky is blue', 'It is raining in Amsterdam', etc. We can explain the reliability of such beliefs by appealing to our perceptual mechanisms, and in most cases, barring universal skepticism, that proves to be enough. However, we also have a lot of beliefs about scenarios and states of affairs that lie beyond actuality. In our day to day lives it often happens that we refer to *what could be the case*. In other words, we have beliefs about what is possible. We would all agree that I could be doing something else instead of writing right now, or that it could be raining tomorrow, or that it must be the case that I am identical to myself. Those beliefs turn out to be essential and tremendously useful, especially beliefs in what is possible. This is because all of us need to engage in various forms of planning throughout our daily lives and planning requires us to consider different alternatives and future possibilities. It is frequently the case that in order to plan or decide on what to do we need to reason about what could happen and consider what the consequences of different possible actions would be. In short, we are engaged in the activity of trying to answer questions of the form '*what would happen if?*', and this is inherently a question about possibility, because it forces us to think of unactualized scenarios that could genuinely come about. Furthermore, we notice various appeals to the realm of possibility in the sciences and in philosophy. Scientists and philosophers alike create thought experiments that essentially involve possible states of affairs and make use of them to justify important claims. We need only think about Chalmers' zombie argument, which relies heavily on the claim that philosophical zombies, i.e. agents with no conscious awareness that are nonetheless functionally identical to agents like us, are genuinely possible (Chalmers, 1995). But whereas in the case of actual truths it is relatively easy and straightforward to explain why we are justified in believing in them, it is a more difficult question how it is that we can explain why we are justified in having beliefs in non-actual and possible

states of affairs. This is the job of *the epistemology of modality*: explaining how we come to gain knowledge of possibilities and how we can justify our beliefs in them. There are two main question that theorists working in this field are concerned with:

(*The Access Question*:) How can we gain epistemic access to the modal space?

(*The Navigation Question*:) How can we epistemically navigate between different types of modality?¹

Modal epistemologies can be distinguished from one another according to multiple theoretical criteria, of which two are most important. First, modal epistemologies can be either empiricist or rationalist (there are also mixed theories but I will ignore them for the sake of brevity). Empiricists usually claim that our justified beliefs in modal facts are grounded in experience and that we acquire modal knowledge through mechanisms that we use in gaining empirical knowledge of the world, like perception, induction, abduction, etc. Rationalists claim that those beliefs are better explained by the existence of some extra-empirical mechanisms that give us epistemic access to modal facts - *a priori* insight, rational intuition, conceptual knowledge, etc. (Vaidya & Prelević, 2023). Naturally, then, and this brings us to the second distinction, most empiricist modal epistemologies are called 'anti-exceptionalist': they explain knowledge of modality by appealing to ordinary cognitive capacities, i.e. capacities that we use when dealing with actual facts. In other words, the claim is that there is nothing special about the things that lead us to modal knowledge, they are capacities or mechanisms that are used towards other, perhaps more mundane, purposes in our day to day lives. Rationalists are usually exceptionalists, in that they claim that we get modal knowledge through the use of non-ordinary capacities or mechanisms. Exceptionalists believe that there is a special capacity or mechanism whose main purpose is providing us with epistemic access to the modal realm. However, this distinction is not always clear-cut, as there are examples of rationalists who are also anti-exceptionalists. One representative of such a view is Amy Thomasson, who is a rationalist that nonetheless locates the origin of modal knowledge in an ordinary understanding of language (Thomasson, 2018). But such examples are rare and most rationalists will opt for a form of exceptionalism. Another point of difference is that most modal empiricists tend to also be naturalists with respect to modal knowledge. Naturalists with respect to modal knowledge claim that we should strive for a *cognitively plausible* epistemology of modality, meaning that it should explain our modal knowledge by appeal to capacities that can be accounted for evolutionarily and developmentally (Goldman, 1994; Jenkins, 2013; Nolan, 2017; Rysiew, 2020) and that *ordinary humans* possess and use in ordinary life. In the words of Williamson:

¹<https://plato.stanford.edu/entries/modality-epistemology/>

A plausible epistemology of possibility should subsume our capacity to discriminate metaphysical possibilities from metaphysical impossibilities under more general cognitive capacities used in ordinary life (Williamson, 2007, p. 136)

While rationalist and exceptionalist approaches provide important and interesting insights into our ways of gaining modal knowledge, I will only focus empiricist anti-exceptionalist modal epistemologies. The reason for this is twofold. For one, I think rationalist and exceptionalist approaches appeal to mysterious capacities that are hard to account for in cognitively plausible ways and they face the challenge of explaining the tension between information about our cognitive development and the exceptionality of those extra-empirical capacities that they appeal to. Furthermore, as Timothy Williamson points out, rationalist accounts usually appeal to capacities or mechanisms for which we have no plausible evolutionary explanation (Williamson, 2007). The second reason for this has to do with one of the aims of the thesis and I will talk about it in detail in the third section of this introduction.

Another important distinction between different epistemologies of modality is connected to *symmetry*. Accounts in the field are primarily interested in explaining how we come to gain knowledge of two types of modal truths: truths about possibilities and truths about necessities. Some theories are *asymmetrical*, in that they give explanatory priority to one type of knowledge over the other. In other words, those accounts claim that we gain knowledge of one type of modal truth first, and then we can explain how we acquire knowledge of the other type of truth through our knowledge of the first type. Asymmetrical theories are either *possibility-first*, when they account for how we gain knowledge of possibilities and then appeal to that in order to explain how we attain necessity knowledge, or *necessity-first*, if they do the reverse. By contrast, *symmetrical* theories do not explain one type of knowledge through the other.

Like any field in philosophy, modal epistemology gives rise to several *problems*. These are conceptual conundrums theories in this field tend to find themselves in and have to solve. The identification of such problems is usually what drives the field forward, by identifying theoretical weaknesses, fatal or otherwise, that then lead to conceptually innovative solutions and further exploration of the logical space of modal epistemology. One such problem will be of interest to me throughout the present dissertation: *the problem of modal epistemic friction*. This is a notion that was coined by Anand Vaidya and Michael Wallner and it concerns the problem of providing epistemic constraints (epistemic friction creators) for the proposed modal guide (A. J. Vaidya & Wallner, 2018). Take imagination as an example. Imagination is a capacity that allows us to envision how things could be, so it is argued that it is one of the guides to modal knowledge. However, our capacity to imagine goes beyond the bounds of possibility, i.e. we sometimes imagine things that are impossible (Byrne, 2007; Fiocco, 2007; Kung, 2014; Priest, 2016; Berto & Schoonen, 2017). The question then is what makes imagination successful in cases in which what the agent imagines is genuinely possible?

1.2 The Problem of Modal Epistemic Friction

The problem of Modal Epistemic Friction (PMEF) is a notion that has been coined by Anand Vaidya and Michael Wallner in their discussion about epistemic constraints in theories about modal knowledge (A. J. Vaidya & Wallner, 2018). As an illustration of their main point, they go through some of the most popular possibility-first epistemologies of modality and point out that each of those theories proposes either a capacity or an inferential schema for the role of modal epistemic guide, but none of them provides epistemic constraints, i.e. *friction creators*, for its candidate. Take conceivability-theory as an example, in the formulation of Chalmers. In short, its main claim is that we are epistemically led from the conceivability of something to its possibility (the position itself is more nuanced than this and Chalmers outlines different types of conceivability, but going into this would take us outside the scope of this thesis). We acquire knowledge of the possibility of P by conceiving of P. However, Vaidya and Wallner raise a critical question: how exactly is it that conceiving of a scenario X can lead us to a reliable judgement concerning the possibility of X? For instance, they take the concrete example of transparent iron. For a subject to conceive of transparent iron she must build a representation R of a scenario S in which transparent iron is present, through a combination of linguistic items, images, or both. Given such an act of conceivability and such a representation, how does the subject get to know that it is possible for transparent iron to exist? Vaidya and Wallner contend that the only way this can happen is if the subject is confident that the constitution of iron does not 'keep R from containing transparent iron'. In other words, the subject can make the jump from conceiving of transparent iron to a possibility judgement as to the existence of transparent iron through knowledge of truths regarding the constitution of iron, i.e. constitutive truths. And since constitutive truths are essentialist truths, the subject is led from the conceivability of transparent iron to the possibility of transparent iron through knowledge about the essence of iron. Thus, it appears that conceivability needs *essentialist constraints* in order to be a genuine modal epistemic guide. After going through the other theories and showing that they are susceptible to the same problem, they conclude that any non-essentialist modal epistemology is, in principle, open to this problem, since any such epistemology will appeal to a capacity or a mechanism that cannot be epistemically reliable unless suitably constrained, i.e. unless there is something that reigns it in by creating modal epistemic friction. For Vaidya and Wallner, those friction creators can only be constitutive truths, i.e. essentialist truths. No matter what we take to be our guide to the possibility of transparent iron, it needs to be constrained by something in order to be a reliable guide, and it cannot be suitably constrained except by appeal to the constitution of, i.e. to the essence, of iron. Thus, epistemic friction is created by essentialist propositions. Not only that, but those essentialist propositions are part of our background knowledge against which 'we can check exercises of conceivability, properly develop a counterfactual supposition, or conclude that something is necessary or possible' (A. J. Vaidya & Wallner, 2018, p. 1921), so the subject

has to know that which creates epistemic friction for the epistemic friction to take place. Consequently, PMEF has two component questions:

- (i) What is it that creates modal epistemic friction (EFC) in an account of modal knowledge?
- (ii) How do we know that which creates modal epistemic friction?

However, a few critical remarks are in order. While Vaidya and Wallner raise an important problem, not all of their conclusions are warranted. First, epistemic friction creators need not be constitutive truths. Embodied imagination-based epistemologies of modality employ recreative or simulationist accounts of imagination. On such accounts, imagination recreates or simulates other cognitive capacities, usually by reactivating motor and perceptual mechanisms that were used by the organism during past experiences (Currie and Ravenscroft, 2002; Nichols and Stich, 2003; Goldman, 2006; Balcerak Jackson, 2018; Gregory, 2020). Consequently, according to simulationists, when we imagine something we activate our perceptual system in the absence of any perceptual stimuli that correspond to the imagined objects. When imagining an apple, we activate perceptual and motor mechanisms that were used in previous interactions with apples. Thus, imagination is architecturally or inherently constrained. However, against Vaidya and Wallner, those constraints are not constitutive truths, since they aren't even the type of thing that can be a truth. They are extra-propositional elements of our cognitive architecture. They are not the type of thing that can be known either, so it could be argued that we have a guide to modal knowledge that is reliable without being constrained by constitutive/essentialist truths and without requiring the subject to know exactly how it is constrained in order for the epistemic friction to take place.

Thus, we are left with an externalist, simplified version of the *PMEF*:

What is it that creates modal epistemic friction (EFC) in an account of modal knowledge?

1.3 Thesis Aims

As I have discussed in the first section, epistemologies of modality usually single out one capacity or (inferential) mechanism as our or one of our guides to modal knowledge. There is one capacity whose status with respect to modal knowledge is contentious: perception. It is a working assumption in modal epistemology that perception cannot be a guide to modal knowledge. It is generally considered that while perception is reliable in giving us information about the actual world, it is silent when it comes to the modal realm (Chalmers, 1996; Hale, 2013; Salmon, 2005). Perception only goes as far as what is actually the case, and if we want to go beyond that, into the space of modality, we need another capacity or inferential schema. However, there have been recent attempts to challenge this assumption, most famously by Margot Strohminger and Bence

Nanay (Strohminger, 2015; Nanay, 2010; Nanay, 2011a; Nanay, 2011b; Nanay, 2014).

The **first aim** of the thesis is to analyze and critique those attempts at identifying perception with one of the routes to modal knowledge. Many theorists talking about this topic have the habit of just quoting Strohminger, Nanay or both and assuming that they are right in claiming that perception is sometimes a guide to modal knowledge, without further argument or closer scrutiny. I will argue that if we look closely at the arguments and try to elaborate a perception-based modal epistemology out of them we run into trouble. In particular, those accounts fall prey to the *PMEF* and collapse into ability-based modal epistemologies.

The **second aim** of this thesis is to offer an alternative, ability-based modal epistemology which explains the cases that perception theorists are interested in. I will argue that in those cases our modal knowledge is better explained by what we know about our own abilities. This ability knowledge is grounded in the process of bodily awareness and in the feeling of bodily ownership.

The first aim of the thesis will be covered by the first three chapters. In the first chapter, I will analyze and critique Margot Strohminger's account of perception as a guide to modal knowledge. I will argue that in refusing the claim that our perceptual episodes have modal content, she makes her account vulnerable to the *PMEF*. This is because then we can ask what exactly it is that leads us from the perception of a tree to the modal knowledge that it is climbable. What leads us from the perception to the modal knowledge is ability knowledge. But then it is ability knowledge that does the heavy lifting, because it is ability knowledge that actually guides us to the belief that the tree is climbable. But then Strohminger's account collapses into an ability-based modal epistemology. This leaves her with a dilemma: she either opts for an ability-based account or she goes for a theory of perception in which our perceptual representations do have modal content, i.e. we perceive modal properties directly. In the subsequent two chapters I will go over theories of perception that make up the second fork of the dilemma, namely affordance theories of perception (both sensorimotor accounts and direct perception accounts) and Bence Nanay's pragmatic approach. For each of the two I will show how we can elaborate a modal epistemology based on them, but I will argue that neither of them works. First, the sensorimotor approach claims that perception requires ability knowledge, so the modal epistemology based on it collapses into an ability-based one. Second, direct perception approaches manage to avoid this issue, but they face the problem of subject relativity: it seems like the information that affordances reveal in direct perception is tied to the specific object that the agent sees in front of them, but the issue is that modal knowledge is usually general in nature. For instance, if we have modal knowledge about the tree in front of us then we know that we can or cannot climb trees in general. Direct perception theorists cannot explain how this happens, how we go from seeing the tree as climbable to the full-blown modal knowledge about the climbability of trees in general. Another issue is that they cannot explain how we maintain the belief that the tree is climbable in the absence seeing the tree. Finally, the modal epistemology

based on Nanay's action-based, pragmatic account falls prey to the *PMEF*. This is because he rejects the claim that our perceptual mechanism reliably tracks modal properties out there in the world. But then we can ask what it is that makes our perception of modal properties successful, in the cases where it gets things right. The answer is, again, ability knowledge, so this approach collapses into an ability-based modal epistemology as well.

Chapter 2

Climbing Trees and Reaching Mugs

2.1 Introduction

Many philosophers working in the epistemology of modality are working with the assumption that perception cannot deliver knowledge of modal truths. Perception can only inform us about what is actually the case and the modal realm goes beyond the epistemic capacities of perception (Chalmers, 1996; Salmon, 2005; Hale, 2013). However, in recent years, there have been several authors, chief among them Margot Strohminger and Bence Nanay, who are attempting to challenge this deep-seated principle. In this chapter, I will present, analyze and critique Margot Strohminger's account of how perception delivers modal knowledge. In the first section I will go over her account, while in the second section I will argue that it fails because it falls prey to the problem of modal epistemic friction (*PMEF*).

2.2 Strohminger's Account

Strohminger (2015) attempts to challenge the assumption by arguing that there are some cases in which we possess visual perception of nonactual metaphysical possibilities. She's operating with a view of perception that conceives of it as a process that is 'psychologically immediate', not including inferences from prior assumptions (Strohminger, 2015). In setting up her case, she starts by analyzing situations in which we report pieces of knowledge by the use of the verb 'see'. We frequently utter statements like 'I saw that the bird had a blue patch' (Strohminger, 2015). Here, it seems that 'see that' can be paraphrased as 'came to know through vision that' and, consequently, it is used to convey that I gained knowledge of something through sight (Strohminger, 2015). The verb 'see' can also be used in different way, as a synonym to 'came to realize',

but such uses don't have anything to do with vision, so we are not interested in them here (Strohming, 2015). Statements containing 'see' in the visual-epistemic sense that Strohming is interested in differ from other 'see'-sentences in that they can be modified by adding 'gaze-directional' prepositions, like 'in the book', or 'through the windshield' (Strohming, 2015). Consequently, we can turn a sentence like 'I saw that there was a protest' into 'I saw on the internet that there was a protest' (Strohming, 2015). However, it should be apparent that this type of situation doesn't capture the epistemic category that Strohming is looking for. This is because this class of 'see'-sentences contains cases in which the knowledge that is reported is indirect (Strohming, 2015). This way of acquiring knowledge of modalities is trivial and uninteresting for what Strohming wants to do. She's not interested in cases where an already established modal fact is communicated in such a way that you can come to know it through vision (Strohming, 2015). She's not interested in cases where we come to see possibilities with the help of third parties, she is interested in cases where we come to see possibilities directly (Strohming, 2015). In other words, we are looking for visual knowledge reports that cannot be modified by gaze-directional prepositions that turn them into statements referring to situations where the seeing subject wasn't a direct witness of the situation (Strohming, 2015). Instead, we are looking at visual knowledge reports that can be modified with 'can'. Sentences like 'I could see that there was a protest' seem to suggest that the seeing subject is a direct eyewitness to the fact that is reported as known (Strohming, 2015).

Strohming then considers cases in which we have knowledge of our own abilities. For instance, we routinely say things like 'I can sit on the chair' or 'I can climb the tree' (Strohming, 2015) and there is a sense in which we see those things (Strohming, 2015). By looking at the tree, I perceive that I can climb it (Strohming, 2015). I see the tree and I judge that I can climb it immediately (Strohming, 2015). There is no reasoning between me seeing it and me judging that I can climb it, i.e. in judging that it is possible for me to climb the tree (Strohming, 2015). In virtue of this happening automatically, without the mediation of other mental capacities or schemes of reasoning, and in virtue of such claims being relatively reliable, we can be said to acquire modal knowledge through perception, at least in this very restricted set of scenarios (Strohming, 2015). Moreover, 'see'-sentences involving personal abilities can be modified with appropriate gaze-directional prepositions. For instance, we can turn 'I see that I can sit on that chair' into 'I see, through my glasses, that I can sit on that chair', or 'I see that I can climb the tree' into 'I see, through the window, that I can climb the tree' (Strohming, 2015). More importantly, I can modify those sentences by the addition of 'can' (Strohming, 2015). In the two example from above, I'm perfectly able to also say the I can see that I can reach the mug, or that I can see that I can climb the tree. Therefore, my perception of those nonactual facts seems to be direct (Strohming, 2015).

Furthermore, Strohming claims that our capacity to perceive possibilities extends even more. Sometimes we can come to know modal facts through perception and inference. For instance, I can see that I can reach the mug and

make the following argument:

P1. I can climb the tree.

P2. If I can climb the tree then it is metaphysically possible that I climb the tree.

C. Therefore, it is metaphysically possible that I climb the tree. (Strohminger, 2015)

We can even extend this to cases that do not involve our abilities. Consider the following argument:

P1. The T-shirt is black.

P2. The T-shirt could have been some color other than black.

C. Possibly, the T-shirt is not black. (Strohminger, 2015)

These are obviously not examples where we acquire modal knowledge directly through perception or even on the basis of perception. Rather, Strohminger makes the point that after performing inferences like the ones above for a sufficient number of times, we get into the habit of getting from the visual perception to the conclusion directly (Strohminger, 2015). In such a case, we don't need the intermediary premise or principle that refers to metaphysical possibility, we can jump from seeing to the possibility judgement directly (Strohminger, 2015). Therefore, once we repeat those arguments enough times, we get to make possibility judgements based directly on visual perception, so we have cases in which sight actually delivers modal knowledge (Strohminger, 2015).

There are a few important clarifications to be made here. First, we should make a distinction between *perceiving* that something is possible or that one is able to do something and judging that this is the case on the basis of perception (Strohminger, 2015). In the first case, the fact that something could obtain is part of the content of the perceptual experience, while in the second case this is not so and the same possible fact is inferred through an additional mechanism or procedure. Strohminger stresses that she does not argue for the claim that modality is included in the content of perceptual experiences related to our abilities. Thus, she is not claiming that when I see the tree, my visual perception has the content that the tree is climbable for me. Her claim does not require this because perceptual knowledge is possible even if the content that is known is not included in the content of our perceptual experiences (Strohminger, 2015). Therefore, Strohminger is merely claiming that my knowing that I can climb a tree is a piece of perceptual knowledge, i.e. knowledge justified by or arrived at through perception. This means that she is not opting for an enactivist or affordance-based theory of perception, since the latter hold that the sensory experiences analyzed by her have modal contents. Second, Strohminger is not claiming that perceptual experiences can *immediately* justify beliefs about non-actual possibilities. Even if you have perceptual justification for holding a belief to be true, it doesn't follow from this that having the corresponding perceptual experience immediately justifies your holding that belief to be true (Strohminger, 2015). Something more is needed.

2.3 Critique

Although this proposal is *prima facie* attractive, I want to argue that it runs into some (potentially) damning issues. First of all, Strohminger is mainly concerned with situations in which we arrive at possibility judgements that have to do with our abilities. Her main examples are 'I can reach the mug' and 'I can climb the tree'. However, as possibility judgements, this class of statements talks about practical possibility, rather than metaphysical possibility, since it refers to a possible scenario whose possibility hinges on specific contextual factors, in this case the subject's ability. Judgements of metaphysical possibility are different, however, since they hold irrespective of situational factors like a person's abilities. There is a difference between saying that the mug is reachable, given my ability, and saying that it is possible to reach the mug. Since modal epistemology is primarily interested in knowledge of metaphysical modality, this raises the following navigation question: If we aren't skeptical about knowledge of metaphysical modalities, how do we navigate from judgements of practical possibility to judgements of metaphysical possibility? Strohminger's account seems to be silent on this. One way for her to reply would be to say that practical modality is included in metaphysical modality, which is to say that everything that is practically possible is also metaphysically possible, simply in virtue of how the modal space is structured. Thus, whenever I know that p is practically possible, I automatically know that p is metaphysically possible. However, this does not necessarily hold and the reason why can be illustrated by a simple example. Imagine that George is a philosophy student and is friends with a law student, John, who studies at the same university. Unbeknownst to George, all law students at their university are good debaters, and that includes John. Now, George knows that John is a law student, but that alone doesn't necessarily put him in a position to know that his friend is a good debater, even though, as a matter of fact, the set of law students at their university is included in the set of good debaters. To know that John is a good debater, it isn't enough to know that he is a law student, even if it's true that all law students are good debaters. To know that John is a good debater, George would either have to watch him debate, or he would have to know that all law students are good debaters. The account she sketched is silent on this, but Strohminger can reply in either of two ways: (1) she can claim that we operate with the tacit assumption that everything that is practically possible is metaphysically possible, or (2) she can claim that we know, explicitly, that everything that is practically possible is also metaphysically possible. I don't think anyone would entertain (2) seriously, it should be obvious that non-experts do not have that kind of modal expertise. As for (1), the claim that we operate with such an assumption is very controversial and needs to be backed up by an argument. Strohminger would need to explain how we form such an assumption or tacit background principle because it is far from clear that we do so.

Another issue that might be raised in response to Strohminger's sketch has to do with her distancing herself from accounts of direct perception. Since she isn't opting for such an account, she doesn't hold that the possibilities in ques-

tion are part of the content of our perceptual experience. This means that arrive at the knowledge of such possibilities only on the basis of perception, with the help of an additional mechanism, be it an inferential procedure or a capacity. But if this is the case then we might ask if it can still be said to be true that we acquire modal knowledge through perception in those cases. Why shouldn't we say that we acquire this knowledge through an inferential procedure or through some other mechanism? Strohminger gives us the beginning of an answer when she talks about situations in which we arrive at possibility judgements through inference, via the principle 'If I can do X , then it is metaphysically possible that I do X '. She goes on to say that even though these examples present us with case in which we acquire knowledge of possibility indirectly, not necessarily through perception, this is only the case initially. After repeating this type of inference multiple times, we internalize it and are able to jump from the perceptual experience of something to the judgement of possibility directly, without the mediating modal principle. However, this is not satisfactory. The question under discussion is not whether there are certain scenarios at all in which we get to possibility judgements on the basis of perception (more or less) directly, but whether there is a class of cases where perception leads us to reliable possibility judgements *in the first place*. Of course there are cases in which perception leads us to modal knowledge, but if the cases in question are just situations in which we jump from perceiving something to a possibility judgement directly because we have learned to do so then we haven't said anything interesting. The question we're interested in is whether we can get to those possibility judgements through perception, more or less directly, in the first place, not by learning to do so. Imagine someone claimed that we sometimes have access to mathematical truths through direct rational intuition and she gave as an example a proof that we go through so many times that we get to the point where we can 'see' the conclusion, a mathematical truth, without going through the intermediary steps of the proof.

Another issue is related to *the problem of modal epistemic friction (PMEF)*. Remember from the introduction that the *PMEF* talks about the fact that epistemologies of modality need to offer constraints, i.e. *modal epistemic friction creators*, for their proposed *modal epistemic guide*, namely for the capacity or mechanism that they claim leads us to modal knowledge. Otherwise, they run the risk of being undermined by the fact that their modal guide requires another type of modal knowledge in order to work as a modal guide. In the Introduction I talked about Vaidya's and Wallner's example of conceivability theory. They argue that what leads us from conceiving a possible state of affairs to the judgement that that state of affairs is possible is essentialist knowledge/information. We know that the conceived transparent iron is possible because transparent iron does not violate the essence or the nature of iron. Otherwise it's hard to see how conceivability could lead us to possibility knowledge. The problem for conceivability theory is that once we accept this, it seems like essentialist information is doing the heavy lifting in the process of modal knowledge acquisition, and thus replaces conceivability as the modal epistemic guide. So the question is: what is it that creates *modal epistemic friction*? Vaidya and Wallner claim

that epistemic friction creators have to be constitutive truths and that we have to provide an account of how we get knowledge of those constitutive truths (Vaidya & Wallner, 2018). However, as I have argued in the Introduction, this is not true. If we look at epistemologies of possibility based on embodied imagination, we see that those are accounts where the epistemic friction is created by *architectural constraints*, namely by motor and perceptual mechanisms that were activated in previous encounters with the thing that is imagined. Those architectural constraints are neither essentialist truths, nor things that we know. But we can still ask for modal epistemic friction creators.

Similarly to Vaidya and Wallner's example of transparent iron, we can ask how exactly it is that perceiving the tree leads us to the possibility judgement that I can climb the tree? Strohminger distances herself from the claim that possibility is part of the content of our perceiving of the tree, so she cannot say that perception leads us to a possibility judgement simply because it contains the climbability of the tree in its content and this can be consequently expressed in a judgement. Therefore, she must either appeal to essentialist truths or to some capacity or inferential schema. If she appeals to essentialist truths, then she must claim that we have epistemic access to the constitution of the tree and that in light of this knowledge we are confident that its essence does not keep the tree from being climbed. However, this undermines her claim that in cases about our abilities it is perception which plays the central role in the process of modal knowledge acquisition. This essentialist picture seems to collapse her account into an essentialist epistemology of modality. It makes the process between the act of perception and the possibility judgement indirect to the point where we can ask whether we have a perception-based modal epistemology anymore, since the heavy lifting in the process of acquiring possibility knowledge is done by the essentialist knowledge. However, she can opt for an inferentialist approach, where we get to the possibility judgement stating that it is possible that I climb the tree through an inference from my perceiving that I can climb the tree, via an instance of the above mentioned modal epistemic principle: 'If I can climb the tree, then it is metaphysically possible that I climb the tree'. I think this is a more promising avenue, but this raises many questions. How do we get to make those inferences? How do we get to know that 'If I can do X , then it is metaphysically possible that I do X ' (since, again, it is this principle that seems to be doing the heavy lifting here)? Can we really say that perception is the guide to modal knowledge here? Why not say that our guide to modal knowledge is an inferential schema that we apply to statements about our abilities? Perhaps Strohminger could reply that this requirement is too stringent. After all, she can say that those inferences mentioned above are internalized, that we perform them automatically and unconsciously. There is a perceptual episode that justifies me to claim that I can climb the tree. We have tacit/background knowledge of the principle 'If I can do X , then it is possible that I do X '. This tacit principle allows us to get from the perceptually informed belief that I can climb the tree to the conclusion that it is possible that I climb the tree. The principle is internalized and works in the background. It is true that sometimes the principle is explicitly stated, but for the most part it works

in the background, more or less unconsciously.

It is important to note that this is a principle of the form ' $P \rightarrow \Diamond P$ ', since it allows us to go from an actuality statement to a possibility statement. At this point we can make an analogy with Kripke's treatment of modal knowledge acquisition, i.e. his 'deduction theory'. In short Kripke states that we come to know necessary facts through the use of an inference schema or pattern, which looks something like this:

- P1. $P \rightarrow \Box P$
- P2. P
- C. $\Box P$

What this says is that we sometimes come to know, by a priori philosophical analysis (perhaps), that if something, call it X , is the case then it must necessarily be the case, and then we discover, a posteriorily, that X is actually the case. Consequently, we conclude that X is necessarily the case. Vaidya and Wallner point out that whereas in the case of conceivability theory the problem of modal epistemic friction arises in relation to the way in which we are led or guided to possibility, in the case of deduction theory the *PMEF* appears in relation to the justification for the first premise, ' $P \rightarrow \Box P$ '. Here we can ask how it is that we come to know principles of this form, or rather how it is that we come to know which instances of $P \rightarrow \Box P$ are true. Again, Vaidya and Wallner offer an essentialist response to this question, a response which they take to show that deduction theory needs essentialist constraints in order to avoid *PMEF*. Here too epistemic friction is created by essentialist propositions. In short, if P is an essentialist truth, then P is (metaphysically) necessary. By knowing that P , we know an essentialist truth, and therefore we know a necessary truth. We get to know the principles that get the argument off the ground through our knowledge of essence.

We can ask the same question of Strohmingner then: how is it that we come to know principles of the form ' $P \rightarrow \Diamond P$ ' and how do we get to know which instances of this form are true? It is tempting to say that a reply that is similar to the essentialist one is available for Strohmingner here. She could respond to the *PMEF* by saying that the modal epistemic friction is created through ability knowledge. When P is a statement about an ability, P is a statement of possibility. Knowledge of ability entails knowledge of possibility and that is how we come to know the true instances of the general principle that gets the argument off the ground. However, I think knowledge of ability does less than what is suggested here. In the case of deduction theory, it is easy to see how essentialist knowledge, if it exists, leads us to knowledge of metaphysical necessity. However, it is not as easy to see how knowledge of one's abilities can lead us to knowledge of metaphysical, as opposed to merely practical, possibility, i.e. the knowledge that we need to secure in order for the argument to work. And even if ability knowledge led us to metaphysical possibility knowledge, the only way to maintain that this type of account is perception-based is to hold that we get to ability knowledge through perception, i.e. by perceiving ourselves as *able to do* such and such. But the perception of our abilities would have to include modality in its content, which is a claim Strohmingner wants to distance herself from. Thus, it seems that Strohmingner either accepts that some perceptual representations include modal content, or she renounces the idea that perception can tell us something about what is possible. Of course, she can also take the route favoured by Barbara Vetter and propose an alternative metaphysics that allows for potentialities, which are modal properties in the world, but even so, it is hard to see how perception could lead us to knowledge of a restricted class

of them without representing them (presumably they couldn't be represented through non-modal content, since they would include some sort of modality by their very nature).

2.4 Conclusion

In this chapter I have presented Margot Strohminger's account of how we can gain modal knowledge through perception. I provided a criticism of this account, on the grounds that it falls prey to the problem of modal epistemic friction (*PMEF*). This is because perception needs ability knowledge in order to lead us to possibility knowledge. Thus, ability knowledge plays a similar role to essentialist information in Vaidya and Wallner's discussion on the *PMEF*. Also, once we accept this, ability knowledge is the one doing the heavy lifting in the process of possibility knowledge acquisition, so Strohminger's account collapses into an ability-based epistemology of modality.

Chapter 3

Pragmatic Representations

3.1 Introduction

In the preceding chapter, I have contended that Strohminger’s theoretical framework is susceptible to the Problem of Modal Epistemic Friction (PMEF). To overcome this predicament, I have argued that there are two potential resolutions: (1) the adoption of a perception theory that facilitates perceptual experiences with modal content or (2) the implementation of an ability-based approach to modal epistemology. In the subsequent chapters, my focus will be on exploring the first alternative, examining two types of perception theory. These include the more conventional perspective put forth by Bence Nanay, which is predicated upon pragmatic representations, and the more unconventional approach of Gibsonian affordance theories. In the present chapter, I aim to delineate Nanay’s perspective and argue that it, too, faces vulnerabilities regarding the PMEF. Specifically, in the initial section, I will expound upon Nanay’s overarching approach to perception and action, which is rooted in pragmatic representations. Following this, in the second section, I aim to elucidate and critically analyze his arguments for the claim that we represent Q-able properties perceptually. On the basis of this approach, I will then try to sketch a modal epistemology where this specific brand of perception-for-action is our modal epistemic guide. Finally, in the final section I will argue that this account also falls prey to the PMEF and that, as far as modal epistemology is concerned, the cases of modal knowledge that we are concerned with are better accounted for by our ability knowledge.

3.2 Action and Pragmatic Representations

For Nanay, in studying the mind, we typically focus on specific mental processes or abilities and conceptualize cognition based on these capacities. This framework allows us to identify two broad strategies for examining the mind (Nanay, 2014). The first strategy targets higher-order, sophisticated cognitive processes

like language and reasoning - traits that are distinctively human (Nanay, 2014). They usually take propositional attitudes or something akin to linguistic entities to be the fundamental building blocks of cognition. This is known as the traditional or cognitivist approach (Nanay, 2014). In contrast, the second strategy focuses on lower-level cognitive processes shared by humans and other animals, those fundamental mental abilities that enable successful daily functioning (Nanay, 2014). Since these processes are evolutionarily older, more fundamental, and more widespread, this approach argues that the study of the mind should prioritize these basic capabilities over the more advanced ones, which are likened to the 'icing on a cake' (Nanay, 2014). When looking at the cake, the base should get more attention than the icing. While the icing enhances the cake's appeal, it does not fundamentally define it (Nanay, 2014).

Nanay supports this second approach, advocating for understanding basic mental processes before tackling the complex ones. Furthermore, Nanay contends that there is no reason to suspect that propositional attitudes, which are the building blocks of language and reasoning, are the fundamental units of cognition. In fact, he claims that in making this assumption we are just projecting the structure of our more conscious mental processes, specifically language, onto the whole of the mind (Nanay, 2014). Abilities like language use and complex reasoning are relatively recent evolutionary acquisitions, as opposed to cognitive capacities that are older and more fundamental for daily tasks, so by projecting the way the first function onto the whole mind we run the risk of obscuring or obfuscating the functioning of the latter, and thus misunderstand how we get to perform more basic tasks (Nanay, 2014). Since the human mind has been selected for the successful performance of daily tasks, it is more appropriate to look at what makes successful action possible in order to discover what the fundamental building blocks of our cognitive apparatus are. Now, Nanay argues that since most actions require perceptual guidance, the building block that we are looking for must be some kind of perceptual state that guides action. It is not a simple perceptual state of the passive observer, it is a perceptual state that is oriented towards performing an action or completing a task and it is responsible for representing the properties that are relevant for the performance of that action or for the completion of that task (Nanay, 2014). Nanay calls those perceptual states pragmatic representations and the properties they represent action properties.

In doing this, Nanay is actually connecting two separate (albeit interconnected) fields: the philosophy of cognition and the philosophy of action. The philosophy of cognition aims at understanding how cognition works, what its functions are, and how it connects to the body and the world, so basically what goes on inside our heads and how it is connected to everything else. The philosophy of action is more interested in understanding what makes action possible. In other words, philosophers of action are interested in explaining what mediates between sensory input (passive perception and beliefs) and behavioral output. It is important to note that what we mean by behavioral output is specifically intentional action. Moving your leg after someone struck your knee with a hammer does not count as action here, since it's just a reflex, an automatic limb

movement. Nanay is connecting the two fields by suggesting that the answer to the central question in philosophy of action (what mediates between sensory input and behavioral output?) is also the answer to the question about what the fundamental building blocks of cognition are. Thus, for Nanay cognition has evolved primarily for the purpose of allowing us to perform actions successfully, and it is first and foremost oriented towards doing things in the world. When it comes to cognition, accurately representing the world and communicating about the world are secondary to action. What is more important is to represent what is necessary for the successful performance of action.

Furthermore, Nanay's account also puts more emphasis on the role of perception in the performance of actions. The cognitivist approach to the mind usually leads to a 'belief-desire' model of action, where the primary role in generating action is played by beliefs and desires (Nanay, 2014). Suppose that I look out the window and I see that it is snowing. This makes me form the belief that it is snowing outside. Since I have the desire to stay warm and comfortable and the belief that the best way of staying like that is to put a winter jacket on, I form an intention to put a winter jacket on. This intention triggers the action of taking the winter jacket and putting it on. However, as Nanay points out, this is not how most of our actions are performed, this accounts only for complex, conscious and uniquely human actions (Nanay, 2014). If we want to account for all kinds of action, we need an alternative account (Nanay, 2014). Nanay suggests that instead of desires and beliefs, the most important role in triggering action is played by pragmatic representations. These are representations that represent action-properties like spatial location, size, weight and, as we shall see, several dispositional properties like climbability, reachability, etc. (Nanay, 2014).

3.3 Do We Represent Objects as Q-able in Perception?

In the last section, we have seen how Nanay accounts for action and how he sees pragmatic representations as being the immediate mental antecedent of action. Pragmatic representations are perceptual states that represent various properties that are necessary for performing actions. But we are more interested in what kind of properties they represent. More specifically, we are interested in seeing whether some pragmatic representations have modal content, i.e., on whether some pragmatic representations represent dispositional properties, like climbability, reachability, etc. Nanay thinks so and presents a series of arguments for this claim. In this section, I will outline and analyze those arguments.

First of all the claim being made is the following:

(PD) There are some dispositional properties that we can perceive objects as having. (Nanay, 2011a)

In other words, there are some properties regarding what objects can do or

what can be done to them that we can perceive objects as having, which is just another way of saying that there are properties about what can be done that we can (and sometimes do) represent perceptually (Nanay, 2011a). For example, we can perceive the mug as reachable, the tree as climbable, etc. There are some properties of the type 'F is Q-able' that we attribute to objects in perception. Those properties are called dispositional because they describe the dispositions of objects. Usually, those properties are considered to be relational in character, and this is because they point to an action that can be performed on an object (that the tree is climbable is just another way of saying that the tree can be climbed). Performing an action on an object depends on both the properties of the object and the capabilities of the agent who is performing the action. According to Nanay, (PD) can be interpreted in two ways, depending on whether we take perception to be extensional or intensional (Nanay, 2011a). If we take perception to be extensional, then the question about what kind of properties we perceive is a question about what properties we track or are sensitive to (Nanay, 2011a). If, however, we take perception to be intensional, then the perception question is about what properties we attribute to objects in perception (Nanay, 2011a). For Nanay, the two interpretations should be kept apart, as they lead to two separate versions of (PD).

On the extensional interpretation, (PD) is saying that our perceptual system tracks or responds to dispositional properties (Nanay, 2011a). That is to say, the claim being made is that dispositional properties are among the properties that causally influence our perceptual system (Nanay, 2011a). If we change a property and our perceptual states are affected by this, then we perceive in the extensional sense (Nanay, 2011a). Thus, as Nanay notes, this interpretation of (PD) has more to do with what properties have an effect on our perception. It is called extensional because if two properties apply to the same set of objects, i.e. if they are coextensive, then if our perception is responsive to the first property this means it is responsive to the second one as well (Nanay, 2011a). On this interpretation, (PD) should be read as (PDE): There are some dispositional properties that our perceptual state responds to (Nanay, 2011a). Depending on the theory of perception we adopt, this can have various consequences. If we adopt a causal theory of perception, then (PDE) implies that dispositional properties are causally efficacious and in talking about it one would then have to engage in all sorts of metaphysical debates (Nanay, 2011a).

On the intensional interpretation, on the other hand, (PD) is simply saying that there are some cases in which we attribute dispositional properties to objects in perception (Nanay, 2011a). In contrast to the extensional case, this is not about properties in the world and whether or not they have an effect on us or are causally efficacious, this is about what our mind attributes to objects, independent of whether or not the attributed properties exist or not (Nanay, 2011a). In other words, the intensional reading of (PD) is about whether or not dispositional properties show up in the content of some of our perceptual states (Nanay, 2011a). On this interpretation, as Nanay notes, (PD) becomes (PDI): There are some dispositional properties that are perceptually represented (Nanay, 2011a). In contrast to (PDE), we can have a case in which two coexten-

sive properties do not have the same effect on the perceptual system, which is to say we can have cases in which our perception is responsive to one property but not to the other (Nanay, 2011a). Furthermore, as Nanay points out, a perceptual state can be caused by a property and at the same time fail to represent that property, to include it in its content. Conversely, a perceptual state can attribute a property *P* to an object that fails to satisfy *P* (Nanay, 2011a).

But note that the question we are interested in, whether or not we represent dispositional properties in perception, is really about (PDI), not (PDE) (Nanay, 2011a). This is because we are interested in finding out whether or not dispositional properties are part of the content of some of our perceptual states, not in finding out whether or not we have metaphysical or causal interactions with dispositional properties located out there in the world (Nanay, 2011a). Thus, we can discard (PDE) and focus on analyzing (PDI) (Nanay, 2011a). The received opinion is that it is false, but Nanay provides three arguments for the (PDI) thesis, two negative and two positive.

The two negative arguments are essentially replies to criticisms of (PDI) and are meant to show that we have no reason to reject the thesis. First, he reviews the most popular argument against (PDI), the so-called Modal Argument. The idea behind this argument is that perception cannot go beyond what is actual, beyond 'objects of simple sight'.

The argument essentially expresses the idea that our perceptual system cannot track or respond to dispositional properties, since they are supposedly located in nearby possible worlds, not in the actual world, and we can only perceive things that are situated in the actual world (Nanay, 2011a). But, as Nanay points out, this is actually an argument against (PDE), not against (PDI), because it argues that our perceptual system cannot be responsive to dispositional properties or that it cannot come into contact with them (Nanay, 2011a). But as we have seen, a scenario in which we represent a dispositional property without having been affected by that property is perfectly consistent with (PDI) (Nanay, 2011a). We can represent all sorts of properties that we are not responsive to, and we can represent objects as having properties they don't have in the actual world (Nanay, 2011a).

Furthermore, the argument doesn't invalidate (PDI) because in order to represent an object as having a property *P* we don't have to represent what makes that object such that it is *P* (Nanay, 2011a). For example, when representing something as being hot we don't need to represent it as having some particular kind and quantity of kinetic energy (Nanay, 2011a). Now, as Nanay points out, a dispositional property is made dispositional by a set of modal facts that are true at nearby possible worlds, but we don't need to perceive or track or be responsive to those facts in order to represent the properties that are made dispositional by them (Nanay, 2011a). For example, I don't have to represent the fact that apples are being eaten in nearby possible worlds in order to perceptually represent the apple before me as edible. In short, as Nanay argues, we don't need to be able to see other possible worlds in order to see objects in the actual world as having dispositional properties (Nanay, 2011a). This means that the argument succeeds in undermining (PDE), but not (PDI) (Nanay, 2011a).

The second important worry with respect to (PDI) is that we cannot maintain (PDI) while rejecting (PDE) without assuming that perception systematically misrepresents the world (Nanay, 2011a). The reason for this is that rejecting (PDE) amounts to saying that our perception is not responsive to dispositional properties, while (PDI) is saying that we represent at least some of those properties, so if we combine them we get the claim that we represent properties that we are not responsive to and that we cannot track (Nanay, 2011a). This would seem to lead to systematic misrepresentation, since there is nothing to make our perception of dispositional properties reliable (Nanay, 2011a). According to Nanay, we should remember that (PDE) is about the causal relationship between our perceptual system and dispositional properties. A rejection of (PDE) would then imply the belief that there is no causal relationship between our perceptual system and dispositional properties, i.e. dispositional bearings have no causal bearing on our perceptual states (Nanay, 2011a). But do the properties we represent perceptually need to cause our perceptual state in order for it to be accurate or correct? For Nanay, there is no reason to suppose that. In fact, many metaphysical frameworks of causation do not require that (Nanay, 2011a).

We can now move on to the positive arguments for the claim that we represent dispositional properties perceptually. Nanay starts with an argument from Susanna Siegel (Siegel, 2006), which aims to show that sortal properties are represented in perception. Suppose we have two experiences, E1 and E2: the experience of seeing a flower before and after taking a course on the typology of flowers. The argument goes like this:

(0) The overall experience of which E1 is a part differs from the overall phenomenology of which E2 is a part.

(1) If the overall experience of which E1 is a part differs from the overall phenomenology of which E2 is a part, then there is a phenomenal difference between the sensory experiences E1 and E2.

(2) If there is a phenomenal difference between the sensory experiences E1 and E2, then E1 and E2 differ in content.

(3) If there is a difference in content between E1 and E2, it is a difference with respect to sortal properties represented in E1 and E2.

(Conclusion) Sortal properties are represented in perception. (Nanay, 2011b)

In other words, the experience of a flower before learning about it is different from the experience of the same flower after having learned about it (Nanay, 2011b). This difference is perceptual in nature and it is a difference in the properties represented in the perceptual states corresponding to the two experiences (Nanay, 2011b). The difference in the represented properties is a difference that pertains to the representation of sortal properties (Nanay, 2011b). Thus, sortal properties are represented perceptually.

Out of all the premises, (3) seems to be the most controversial and Siegel considers a potential challenge to it (Nanay, 2011b). It could be argued that if there is a difference in content between E1 and E2, then that difference is due to the Gestalt-properties represented in the two experiences (Nanay, 2011b). However, she gives an example where this Gestalt-property approach breaks down. Suppose I look at a facial gesture that expresses doubt before and after I know that it expresses doubt (Nanay, 2011b). Now, it seems implausible to claim that the difference between the two experiences is a difference between the Gestalt-properties that are being represented (shape, color-properties, etc.), so it must be a difference in sortal properties, again (Nanay, 2011b). The argument goes like this:

- (a) E1 and E2 represent the same Gestalt-properties.
- (b) Still, there is a phenomenal difference between E1 and E2.
- (c) The phenomenal difference between E1 and E2 must be a representational difference (this follows from step (2) above).
- (Conclusion) Hence, E1 and E2 must represent different sortal properties. (Nanay, 2011b)

However, Nanay thinks that this breaks down if we distinguish between two ways of attending to properties. We can perceive things while attending to them, but we can also perceive things while not attending to them (Nanay, 2011b). Thus, we can represent Gestalt-properties pre-attentively or post-attentively (Nanay, 2011b). Thus, we have two claims: (1) E1 and E2 represent the same Gestalt-properties pre-attentively and (2) E1 and E2 represent the same Gestalt-properties post-attentively (Nanay, 2011b). As Nanay points out, in order to work, the argument needs (2), but there are reasons to doubt the truth of (2). First, it is highly unlikely that after learning that the facial gesture expresses doubt I will attend to the same features that I attended to the first time (Nanay, 2011b). As Nanay observes, I will be attentive to certain details and try to see how it expresses doubt. The difference can be explained by a difference in the Gestalt-properties that are being represented in E1 and E2 (Nanay, 2011b). More generally, the difference in these instances seems to be given by a difference in attention. In both cases, I learn something that modifies my attention, which in turn causes a change in which properties I am attending to perceptually (Nanay, 2011b). Thus, Siegel's argument cannot rule out an explanation in which the difference between E1 and E2 is a difference between the properties represented post-attentively in E1 and the ones represented post-attentively in E2, so there is no reason to believe that the difference lies in the sortal properties being perceived in the two experiences (Nanay, 2011b). In other words, (a) isn't necessarily true and this also affects the original argument, because it blocks the jump from (2) to (3) (Nanay, 2011b). It seems perfectly reasonable to assume that after learning something about the flower in question we will post-attentively represent in E2 Gestalt-properties that we weren't aware of in E1, like overall shape, color of the petals etc. (Nanay, 2011b).

Nanay proposes that we change the argument so that it is about action-properties instead of sortal properties. Suppose, again, that we have two different experiences, $E1^*$ and $E2^*$. $E1^*$ is the experience of representing a mug as reachable for me and $E2^*$ is the experience of representing a mug as reachable for my twin sister. We can use $E1^*$ and $E2^*$ to turn Siegel's argument about sortal properties into an argument about action-properties (Nanay, 2011b), in the following way:

(0*) The overall experience of which $E1^*$ is a part differs from the overall phenomenology of which $E2^*$ is a part.

(1*) If the overall experience of which $E1^*$ is a part differs from the overall phenomenology of which $E2^*$ is a part, then there is a phenomenal difference between the sensory experiences $E1^*$ and $E2^*$.

(2*) If there is a phenomenal difference between the sensory experiences $E1^*$ and $E2^*$, then $E1^*$ and $E2^*$ differ in content.

(3*) If there is a difference in content between $E1^*$ and $E2^*$, it is a difference with respect to the action-property represented in $E1^*$ and $E2^*$.

(Conclusion) Action-properties are represented in perception. (Nanay, 2011b)

What blocked Siegel's argument was essentially the idea that it makes sense to attend to other Gestalt-properties (not only to other sortal properties) after learning something new about the object that is being experienced (Nanay, 2011b). As Nanay argues, this is the reason why the jump from (2) to (3) is not legitimate. Is it the same with the jump between (2*) and (3*)? Nanay doesn't think so, and the reason for this has to do with the fact that in $E1^*$ and $E2^*$ perception is geared towards action. In the case of experiencing the flower, it made sense for the perceptual agent to change the properties they are attending to (Nanay, 2011b). After all, learning something about an object makes you pay attention to and see things in that object that you were unaware of before (Nanay, 2011b). But $E1^*$ and $E2^*$ are different because they are about the climbability of a tree for two different people, they do not involve an epistemic change in the relationship between the agent and the tree (Nanay, 2011b). After all, as Nanay notes, we have the same climbable tree in front of us and we can assume that we know the same things about it. If we suppose that we have similar climbing skills, then it seems highly unlikely that there could be properties that I attend to when I experience the mug as reachable for me and that I don't attend to when experiencing the mug as reachable for my twin sister (Nanay, 2011b). Nonetheless, Nanay claims, there seems to be a difference in content between the two experiences. Obviously, I experience my own actions differently than I experience my twin sister's actions. And since there is no difference in the non-action properties that we represent the mug as having, the

difference must be in the action-properties that are being represented (Nanay, 2011b). The argument looks like this:

- (a*) E1* and E2* represent the same Gestalt-properties both pre-attentively and post-attentively.
- (b*) Still, there is a phenomenal difference between E1* and E2*.
- (c*) The phenomenal difference between E1* and E2* must be a representational difference (this follows from step (2*) above).
- (Conclusion) Hence, E1* and E2* must represent different action-properties (Nanay, 2011b)

The second positive argument is, again, about perception in action and aims to show that in order to perform an action Q on an object, we have to represent the object as Q-able, and that we do so perceptually (Nanay, 2010). Nanay begins by defining an action as Q-able if it is not impossible for a specific agent. In contrast to Gibson, Nanay claims that a Q-able property of an object expresses not what we should do with the object, but, more modestly, what we can do with it (Nanay, 2010). Thus, for Nanay, the property of being Q-able is not the same as the property of affording action Q. Finally, the action of Q-ing is usually a sufficiently mundane and concrete goal-directed action (we need to exclude the case of very complex actions like winning a war) (Nanay, 2010).

What happens when we perform an action? As we have seen earlier, a necessary condition for action is an antecedent mental state that triggers the action (Nanay, 2010). In this argument, Nanay talks about 'executive representations', a notion coined by Kent Bach. He points out that when we perform an action we represent what is being done and what will be done next (Nanay, 2010).

It should be clarified that the representation in question does not have to be conscious (Nanay, 2010). In fact, it usually happens unconsciously. Thus, in the course of an action I need to represent what I will do next, but not as a prediction, i.e. not as something that is most probable, but as something that will happen (Nanay, 2010). If I am reaching for a mug, I have to represent my hands as if reaching the cup will be the result of what I do with them (Nanay, 2010). In other words, I have to represent the state of affairs that will lead to me reaching the cup (Nanay, 2010). I need to represent the specific hand movement that will lead to me grabbing the cup, not a prediction about my uncertain hand movements (Nanay, 2010). But, as Nanay points out, it is hard to see how I could do this if I didn't represent the cup as grabable.

To put it more clearly, in grabbing the cup, I have an action, moving my hands in a specific way such that I attain the state of affairs in which I hold the cup in my hand, and the goal of that action, holding the cup (Nanay, 2010). Many philosophers have pointed out that performing an action requires the representation of the goal toward which that action is directed, in this case the hand holding the cup (Nanay, 2010). However, beyond this, according to those philosophers, the goal is not represented as a prediction or a wish, but as something that I am in a position to achieve (Nanay, 2010). As Nanay points

out, this means that we also have to represent a way of getting to that goal, in addition to the goal itself. Thus, when performing an action we represent the goal together with the way in which that goal will be brought about (Nanay, 2010). However, it's hard to see how one could do that without representing that goal as attainable (Nanay, 2010). Therefore, I need to represent that cup as within reach to perform the action, which means that I have to represent the cup as reachable (Nanay, 2010).

But, as Nanay argues, that does not yet show that we perceive the cup as reachable. After all, there are two ways in which the cup can be represented as reachable: perceptual or non-perceptual (Nanay, 2010). Let us suppose that the representation of Q-ability is non-perceptual. That means that it is a belief that such and such object is Q-able, a belief that arises through some kind of inference out of some perceptual states (Nanay, 2010). We see an apple that is red and circular and because I know that red and circular things are edible I infer, unconsciously, that the apple is edible (Nanay, 2010). As Nanay points out, this implies that the representation of an object as Q-able depends on other beliefs and on background information. For example, if I believed the object were situated behind a glass wall I wouldn't represent it as reachable (Nanay, 2010). My belief in the object's reachability is sensitive to my other beliefs, in contrast to my perceptions, which seem to have at least a higher degree of independence from my background information (Nanay, 2010). However, this is not consistent with our experience of representing Q-ability before performing actions (Nanay, 2010). Nanay provides an example that is meant to show that representations of Q-ability enjoy the same kind of independence that perceptual states do. Suppose that I am standing behind a glass wall and that I am aware of this. If you throw a ball at me I will reach forward to catch it, even though I know that I cannot catch the ball. This does not seem to be a case of mere bodily movement, such as a reflex. Unlike in the case of a reflex-like movement, I can refrain from reaching for the ball if I want to, the movement that I make doesn't surprise me (it's intentional), and the way I move varies as my representation of what is happening varies, so I am indeed performing an action. Since reaching for the ball requires representing the ball as reachable, I am representing the ball as reachable. But if this representation is a belief or a non-perceptual mental state, then it should be sensitive to my other beliefs and to the other information I hold. But since I know that the ball is impossible to reach and I still reach for it, it must be that my representation of its reachability is insensitive to my other beliefs, so it must be a perceptual representation (Nanay, 2010). Therefore, we perceive some objects as being Q-able when performing actions.

To conclude, for Nanay perception is oriented towards action and performing actions requires a set of perceptual states that represent the properties which are necessary for the performance of those actions. He calls those perceptual states pragmatic representations, and they represent action-properties, which encode the information we need to go through with actions. As we have seen above, the representation of dispositional properties is necessary for action, so pragmatic representation represent dispositional properties. Since they are perceptual, it can be said that we perceive dispositional properties whenever we do things in

the world, whenever we perform actions. This could give way to a sketch about how we come to know possibilities. We are active organisms and we frequently engage in goal-directed action. Whenever we do so, we use pragmatic representations, which are different from normal perceptual states in that they are geared toward action, and represent certain action relevant properties. Some of those properties are dispositional. Most of the time, pragmatic representations are unconscious, but once we repeat an action for a sufficient number of times the contents of the pragmatic representations necessary for its performance become conscious and we become aware that the objects engaged in the action have dispositional properties, which lead to us gaining modal knowledge about those objects.

3.4 PMEF

In the last section we have seen that Nanay has managed to show that we perceive dispositional properties. Since dispositional properties are modal properties, he has managed to show that we perceive modal properties. To be more precise, he has argued that we attribute dispositional properties to objects in perception. But the question we are interested in is whether those acts of perception can lead us to modal knowledge. On the face of it, as we have seen at the end of last section, it is possible for perception to lead us to knowledge of possibility. We engage in action all the time, and since action requires the (mostly unconscious) representation of modal properties, then we are bound to become aware of those dispositional properties by repeating certain actions for a sufficient number of times.

However, there is a problem stemming from Nanay's acceptance of PDI and rejection of PDE. By rejecting PDE, Nanay accepts the idea that perception is not sensitive to or does not track dispositional or modal properties. However, he accepts the claim that our mind attributes modal properties to objects in perception. But then the question is: in the cases in which the attribution of a dispositional property is successful, what makes it so, if our perceptual system does not track and is not sensitive to modal properties? This is a variant of the PMEF, since we are essentially asking what constrains perception in successful cases, in cases in which we are lead to modal knowledge. In other words, what is it that creates modal epistemic friction in the case of modal perception for action?

Again, since we are talking about cases of perception that is action-oriented, it seems like the most plausible answer is knowledge of abilities. Our pragmatic representations are usually right in their attribution of dispositional properties when we know that we can do what we set out to do. Of course, every goal-directed action involves a belief in our ability to reach that goal, but I am not talking about mere belief here, I am talking about full-blown knowledge of ability, a belief that is true and justified. And this should come as no surprise. After all, for most actions, excluding cases of luck, it takes a sufficiently large number of repetitions to perform them successfully and it frequently happens

that we are wrong in representing an object as Q-able because we are not good enough to be able to engage in Q-ing successfully. This means that we get it wrong many times when we represent something as Q-able without having the knowledge that we are able to Q. However, once we gain the knowledge that we can Q, we usually get things right in our attributions of the dispositional property of 'being Q-able' across contexts, and our perception of Q-ability is accurate. To give a concrete example, the first time I engage in climbing a tree, if Nanay is right, I represent the tree as climbable, but not being trained in climbing, I fail. Since I failed, I was wrong in my attribution of climbability and I fell short of knowledge. However, after having trained and after having become confident and knowledgeable in my ability to climb trees, my representation of the tree as climbable is accurate, true and also justified, because I know what I can do with respect to the tree. Thus, in the first scenario I lack modal knowledge, while in the second I don't. If perception were the guide to modal knowledge, then the two scenarios would have to differ in perceptual states. But the difference between the two scenarios lies in my ability knowledge, not in my perception! So what leads us to the knowledge that it is possible to climb the tree must be the ability knowledge I have.

3.5 Conclusion

In this chapter I presented Bence Nanay's account of perception. I then showed how we can elaborate an epistemology of modality based on it and I argued that it ultimately fails. This is because Nanay rejects the idea that our perceptual system reliably tracks modal properties out there in the world and he only accepts the claim that modal properties are merely represented in perception. This means that perception sometimes gets it wrong with respect to modal properties. But then we can ask: what is it that creates modal epistemic friction in the cases in which perception is successful in accurately representing modal properties? This is a version of the *PMEF*.

Chapter 4

Affordance Perception and Possibility

4.1 Introduction

In the previous section, I have explored Bence Nanay's approach to perception, based on pragmatic representations, and I have sketched a modal epistemology based on this account. Then I proceeded to criticize that sketch on the basis that it falls prey to the PMEF and collapses into an ability-based modal epistemology because of that. In this chapter I will explore affordance theories of perception and their prospects of helping us sketch a modest and empiricist modal epistemology. In the first section, I will present the original account of affordance perception, found in Gibsonian ecological psychology. In the second section, I will outline subsequent accounts, that aim to clarify and precisify Gibson's approach, and I will focus on Chemero's theory of affordance perception, which sees affordances as relations between the environment and the perceiving organism's abilities. Finally, in the last section I will attempt to sketch a modal epistemology based on this general way of thinking about perception and I will argue that it too is vulnerable to a version of the PMEF. I will discuss an alternative modal epistemology, based on the affordance perception approach pioneered by McClelland, which seems to avoid the problems that the other theories face, but it will be argued that even though it does that, it collapses into an ability-based modal epistemology.

4.2 Affordances

An affordance refers to something that is afforded by the environment surrounding an animal. The term was coined by James Gibson in his book on ecological psychology, where he claims that we don't just perceive objects in the environment, we also perceive what the environments affords. In other words, we

perceive possibilities for action, we perceive what can be done to objects in the world, in addition to perceiving those objects (Gibson, 2014). Gibson’s radical suggestion is that the basic properties of something in the environment constitute its affordances, so when we perceive those basic properties we also perceive what the object affords (Gibson, 2014). For instance, if we take a surface that is horizontal, flat and rigid, at least to a sufficient degree, then that surface affords sitting, it is sit-on-able (Gibson, 2014). It is also walk-on-able and run-over-able. Importantly, it is sit-on-able for organisms like us, and it is not so for other kinds of animals (Gibson, 2014). Water creatures do not relate in the same way to a terrestrial surface, they cannot sit on it, walk on it, or run over it, because their body would not allow such a thing. Thus, whereas physical object properties are measured and considered independently of anything else, affordances are measured relative to the animal that is afforded an action (Gibson, 2014). Like the action-properties of the previous chapter, affordances are relational in nature and they cannot be reduced to the object that does the affording (Gibson, 2014). Affordances can change if we talk about different animals, they are not simply out there in the world (Gibson, 2014). Something that is afforded to us may not be afforded to a fish or a bird, and vice versa. This is because affordances depend on our bodily makeup and capabilities (Gibson, 2014).

Thus, it seems like every type of animal or organism is sensitive to a specific set of affordances (Gibson, 2014). In other words, it seems like every type of animal or organism is afforded a specific set of things by the environment (Gibson, 2014). Gibson calls this ‘a niche in the environment’. Every species is drawn to a subset of the environment that it uses in order to live (Gibson, 2014). Different species are drawn to different subsets of the environment (Gibson, 2014). For instance, we are not drawn to underwater environments, as opposed to fish, but to terrestrial spaces, and inside the terrestrial environment, to specific kinds of spaces and not to others. Gibson suggests that this niche is constituted by what the environment affords each species. In a sense, an organism’s body and bodily abilities coupled with the physical properties of the environment causes it to be drawn to certain objects or subsets of the environment (Gibson, 2014). In other words, the organism and the niche are complementary (Gibson, 2014). This does not mean that affordances are subjective (Gibson, 2014). The environment of planet Earth could be said to afford animal life even in a possible world where animal life hasn’t developed on Earth. As Gibson notes, the point is rather that an affordance is not open to any kind of animal, because it cannot be taken advantage of by any animal. The flat, horizontal and solid terrestrial surface affords sitting, objectively, but it doesn’t afford sitting for any kind of animal, because not every species is biped and able to sit up straight, to have an upright posture, which is necessary for sitting (Gibson, 2014). This is why Gibson notes that an affordance is neither subjective nor objective, and it is both. It doesn’t really make sense to wonder if an affordance is a property of the object or a property of the subjective human mind (Gibson, 2014). It is rather a property that arises out of the interaction between organism and environment (Gibson, 2014). This becomes quite obvious when we consider the

relationship between the human species and the environment (Gibson, 2014). The environment didn't use to afford as many things as it does now. Over the course of our development and history, we have changed the environment in a way that has created countless new affordances (Gibson, 2014).

4.3 Recent Approaches

Gibson also claims that we perceive affordances directly. This amounts to the radical claim that we literally see what the environment affords us. We see the cup as grabbable, the ball as throwable, the seat as sit-on-able. We do so not only in the context of performing actions, but always, even when we are passive. His theory of perception is in the family of direct theories of perception. This family of theories is to be distinguished from theories where perception is conceived of as indirect. In those theories, perception is usually inferential: we receive sensory information through our receptors, which leads to a sensation and the brain performs an inference on that sensation, creating a meaningful perception (Chemero, 2003). In theories where perception is direct, 'meaning' is out there in the world, it doesn't have to be inferred, we see it directly (Chemero, 2003). In other words, as Chemero notes, we don't need the mind to give meaning to the sensory information we receive, the world is already meaningfully structured. We do not need to infer what can be done, what is afforded to us, the affordances are already out there in the world and we perceive them directly (Chemero, 2003). Thus, affordance perception implies that affordances are part of our ontology, that our ontology should also include affordances (Chemero, 2003). However, as Chemero points out, Gibson is not very clear on what kind of worldly thing affordances are, and clarifying this is crucial for introducing affordances in our ontology. This is made more difficult by some of the ways in which Gibson talks about affordances (Chemero, 2003). In particular, Chemero singles out this passage:

An important fact about the affordances of the environment is that they are in a sense objective, real, and physical, unlike values and meanings, which are often supposed to be subjective, phenomenal, and mental. But, actually, an affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective-objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical, yet neither. An affordance points both ways, to the environment and to the observer. (Gibson, 2014, p. 129)

This makes affordances seem like mysterious entities whose existence no serious scientist would accept (Chemero, 2003). There have been multiple attempts to fix this and offer clearer and more precise accounts of what affordances are (Chemero, 2003). As Chemero points out, most of them take them to be animal-relative properties of the environment, so properties of the environment that are

relevant to a particular type of animal. It is important to note that there are two components to this debate: (1) what kind of animal-relative properties affordances are and (2) what it is about animals that affordances are relative to (Chemero, 2003).

Concerning (1), there are broadly two general approaches: selectionist and dispositionalist (Chemero, 2003). For selectionists, affordances constitute resources in the environment that are exploitable by certain animals (Chemero, 2003). As Chemero observes, on the selectionist view then, affordances are just like other things in the world that we find and use for survival. Furthermore, in terms of evolution and natural selection, the animals that are selected for survival are the animals that have evolved their perceptual system to be able to perceive affordances (Chemero, 2003). Thus, according to selectionists, affordances, just like other resources, are a source of selection pressure (Chemero, 2003). Developing the ability to perceive them is an evolutionary advantage and, consequently, affordances are those resources that some members of a species develop the ability to perceive (Chemero, 2003). As Chemero points out, those are the members that survive. They are resources that essentially drive species evolution forward by creating pressure for certain animals to develop the biological tools necessary for perceiving those resources, those elements of the environment that can be exploited for the reproduction of life (Chemero, 2003). Finally, by treating affordances as resources that enter the cycle of evolution and natural selection, the selectionist view sees affordances as things that exist independently of animals (Chemero, 2003).

The dispositional approach, on the other hand, contends that affordances are dispositional properties of the environment, potential properties that are actualizable with the help of animal intervention or behaviour (Chemero, 2003). As we have seen in the previous chapter, the dispositional properties of an object refer to its tendencies, tendencies that might or might not become actualized, depending on certain contextual factors. A branch is breakable, but only in very specific circumstances, namely those in which it is bent to a sufficient degree (Chemero, 2003). Something is breakable only when there are possible situations in which it breaks. This implies that the manifestation of dispositional properties is tied to certain actualizing circumstances and, more specifically, on someone or something performing an action on the objects that have those dispositional properties (Chemero, 2003). In particular, the actualization of affordances depends on the behaviour of certain animals (Chemero, 2003). A tree is climbable only if there are animals who have a body that allows for climbing, i.e. animals that are capable of climbing it (Chemero, 2003). This means that affordances depend on the presence of certain types of organisms or animals and that they are dispositional properties that are complemented by the properties of animals (Chemero, 2003). They cannot exist as affordances without the existence of those complementary properties (Chemero, 2003). Affordances have no independent existence.

Moving on to (2), there are a couple of suggestions as to what affordances are relative to. One suggestion is that affordances are relative to effectivities, which are animal properties that allow those animals make use of affordances

(Chemero, 2003). For instance, if we go back to the tree example, a relevant effectivity here is whatever bodily ability that allows me to climb the tree, i.e. to use the climbability affordance. Another way of expressing this idea is to say that affordances are relative to certain abilities that allow certain animals to make use of them. Another suggestion found in the literature is that affordances are relative to body scale, an idea that is usually cashed out in terms of measurements that show a correlation between body scale and certain environment features that are relevant to specific affordances (Chemero, 2003).

However, Chemero offers an alternative view. He criticises the idea that affordances are properties of the environment. He claims that we should make a distinction between properties and features. Perceiving that my piano is out of tune involves multiple things: (1) perceiving a particular object, the piano, (2) knowing that it is my piano, (3) knowing what it means to be out of tune and (4) perceiving this particular object, my piano, as having this specific property, being out of tune (Chemero, 2003). In contrast, as Chemero argues, seeing that it is snowing, I do not need to attribute a property to any object or entity. I just need to be aware of the feature of a situation, snowyness, which falls short of the status of property (Chemero, 2003). For Chemero, affordances are more like features rather than properties. When we perceive that something is afforded to us in the environment we don't really attribute a property to a specific object, we just become aware of a feature of that specific situation (Chemero, 2003). Furthermore, they are not features of the environment, as Chemero points out, rather they are features of the animal-environment coupling or system. In other words, they are relations between animals and the environment, or rather relations between animals and certain features of the environment (Chemero, 2003). When perceiving an affordance, I am perceiving a relation between me and certain features present in a subset of the environment (Chemero, 2003). Thus, for Chemero an affordance is a relation between organism and certain features of the environment that has the role of affording a certain kind of behaviour. Chemero defines affordances in the following way:

Affords- ϕ (environment, organism), where ϕ is a behavior. (Chemero, 2003)

But what exactly is it in the organism that the affordance relates to features in the environment? For Chemero, the answer lies in abilities. Since affordances afford behaviour, they must relate the environment with my abilities. When I say that the tree affords climbing, I am perceiving a relation between the tree, or the features of the tree, and my ability to climb.

4.4 Circularity and PMEF

Thus, according to affordance theories of perception, we perceive what the environment affords. This can mean several things, depending on the theory one opts for: (1) we perceive an animal-relative property of the environment or (2)

we perceive a relation between the environment and our ability. However, no matter how we cash out affordances, all of those theories seem to have one claim in common, namely that we perceive what we can do. We have two types of perception, namely unconscious perception, which includes perceptual states in which we see things without attending to them, and conscious perception, which includes perceptual states in which we see things and attend to them. Since we are interested in modal epistemology here, we should attempt to sketch an approach to modal knowledge based on these theories of perception.

According to affordance theories, we literally see what we can do. We perceive the door as something that can be opened, we perceive the mug as something that can be grabbed. When we engage in conscious perception, we are also aware that the door is open-able, that the cup is grab-able, etc. Thus, we are in a position to know what we can do, simply by perceiving. Since the knowledge of ability is the knowledge of genuine possibilities for action, this amounts to modal knowledge. One important idea in affordance perception accounts is that, as McClelland (2019) points out, perceiving affordances provides us with information about the environment and ourselves in equal measure (McClelland, 2019). By perceiving affordances, we also gain information about what we are able to do. For instance, by perceiving a mug I get the information that I am in a position to grab it, that I am *able* to grab it. This is so because, as we have seen in the previous section, no matter how we account for them, affordances are about both the perceiving organism and the environment being perceived. Different types of organisms are afforded different activities by different environments. By perceiving affordances, we get information about the environment, but also about our abilities with respect to that environment. In conclusion, since ability knowledge is a form of modal knowledge, we get modal knowledge through perception.

4.4.1 Two Options

There are two main approaches to this, or rather two ways of cashing this out. One of them sees perception and action as being intimately linked. An example of this is the action-space model, which rests on the claim that conscious perception requires us to have some form of knowledge of where we stand in the action space. The action space is a space for possibilities of action and more specifically a space containing the possibilities of accomplishing one's goals (Ward et al., 2011; Grush, 2007). Thus, we get modal knowledge by perceiving what we can do. But there are a few issues with this proposal. First of all, there is a circularity issue. It seems like by saying that we perceive what is possible for action we are already presupposing modal knowledge, at least in the case of conscious perception. In other words, it seems like knowledge of ability is already baked into perception and by invoking perception we are already invoking modal knowledge. There is no explanation of how we gain modal knowledge, just the presupposition that it is there all along, in virtue of perception. But how do we get to know the ability knowledge that is involved in conscious perception? Again, this is just another version of the PMEF, which asks how perception is

constrained when it guides us to modal knowledge. By requiring us to know what we can do in order to perceive, this account does not explain how we get to modal knowledge, since some form of modal knowledge is already always there. In other words, by making ability knowledge, which is a species of modal knowledge, a pre-condition for possibility knowledge, this account presupposes what it needs to explain. Consequently, it leaves us with an unanswered question: how do we acquire the modal knowledge needed for perception?

However, there is another, more promising way of cashing this out. Affordance perception does not necessarily require ability knowledge for perception. It is possible to claim that we simply perceive what the environment affords us, directly, without prior knowledge of our abilities and without knowledge of our location in the action space. For such an account, the order and role of ability knowledge vis-a-vis perception is reversed, in comparison to the action-space accounts. We do not need to know what we can do prior to perceiving, we learn what we can do by perceiving affordances. Since ability knowledge is a form of modal knowledge, we acquire knowledge of possibility through perception. This type of account does not suffer from a circularity problem, since ability knowledge is something we get to after and through perception. Furthermore, it also doesn't fall prey to the PMEF, because we see possibilities directly and there is no need for any constraints in the process where perception guides us to possibility knowledge. There is no guiding to be done, since perception simply presents us with possibilities for action. However, there are a few reasons to doubt that this type of affordance perception can guide us to ability knowledge. As will be discussed in the next chapter, when we are talking about ability knowledge we are interested in knowledge of robust abilities, not simple abilities (Vetter, 2024). Simple abilities refer to action possibilities that are in principle available to us (Maier, 2022). For instance, I have the simple ability to play basketball, since I could, in principle, learn how to throw hoops. But that doesn't mean I can perform while playing basketball. That requires the presence of a robust ability, which entails a level of control over the action of playing basketball. Knowledge of simple abilities does count as possibility knowledge, but it is harder to acquire and it comes later in our development (Vetter, 2024). We seem to acquire knowledge of robust abilities earlier than we acquire knowledge of simple abilities, so an affordance account must explain how we acquire the former (Vetter, 2024). But affordances do not tell us anything about what action we can control and initiate, they tell us something about what actions are open for us, irrespective of our control and performance levels. And even if affordances tell us something about what we are able to do robustly, they are object-dependent, because they only tell us something about what can be done to a particular object, they do not give information about what the perceiving agent can do in general (Vetter, 2024). For instance, an affordance can tell me that I can climb a particular tree in front of me, but it doesn't tell me that I can climb trees in general. However, ability knowledge is usually present in this general form: when I have knowledge of my ability to climb, I know that I can climb trees in general, my knowledge isn't dependent on a particular tree only. Furthermore, we usually form the conscious belief that we can climb the tree in

front of us only when we already have the knowledge that we can climb trees in general. Finally, we have ability knowledge in the absence of perceiving the objects that we can act on, but on affordance accounts our knowledge of what can be done seems to be tied to perceiving the objects in front of us. But then it seems hard to see how an affordance account can explain the generality of ability knowledge.

4.4.2 A way out?

An alternative proposal comes from McClelland, who suggests that affordances in the environment solicit a behavioural response, in that they trigger the activation of certain motor mechanisms, sometimes without our awareness or intention to perform an action (McClelland, 2019). Based on this, one could claim that we are not necessarily aware of what can be done when we perceive affordances, they just call us into action without us being conscious of it. We become aware of what we can do after engaging in certain actions for a sufficient number of times and learning through those experiences. By learning about our abilities, we learn about possibilities for action, we learn about a subset of what is possible. This opens up an interesting avenue in the epistemology of ability, since affordances are not something that we are necessarily conscious of perceiving, i.e. that we see consciously, but something that calls us into action. More precisely, they are something that activates our motor systems, without us necessarily having any intention to engage in any action. This means that affordances sometimes push us forward into engaging in exploratory and unintentional behaviors that can lead to learning about our abilities and thus gaining ability knowledge. As will be discussed in the next chapter, this exploration can trigger interest in certain types of activities and make us particularly attentive to the means of performing those activities (Fridland, 2013). This phenomenon is called *intrapersonal imitation* and it is a part of our normal development. By engaging in intrapersonal imitation, children learn to turn their attention to their own actions and 'imitate' them through repetition (Fridland, 2013). In trying to imitate their own actions, children get attuned to the particular ways in which those actions can be performed and learn to control them (Fridland, 2013). By learning to control them, they form abilities and ability knowledge. The ability knowledge comes about through the development of certain kinds of representations about one's own actions. As will be discussed at length in the next chapter, ability knowledge requires the development of a representation of what the agent can do, in the absence of the context or environment that first caused the action to be learned and controlled. For instance, take a child that is standing in front of a balloon. Let's say the child trips and hits the balloon, causing it to move forward, hit the wall and bounce off it. If the child then engages in intrapersonal imitation and learns the ways in which they can move their body so as to make the balloon bounce off the wall repeatedly, then by repeating and exploring the action they learn to control it and in time develop a representation of what they can do to any sufficiently similar balloon, in any room that has walls. It is only then that the agent can be said to know they have

the ability to play with a balloon. This is because ability knowledge requires a sense of initiation and control. I know I can do something because I can initiate that action at will and I can control it reasonably well. But in order to do that I need to have a representation of the sequence of moves necessary for that action and I need that sequence to not be dependent on a specific environment or context. For initiation at will the agent needs to be able to do the action in multiple contexts and for control the agent needs to have the information about how to perform the action. But notice that even though this process has been kicked off by affordance perception, what leads us to ability knowledge is not affordance perception itself, but intrapersonal imitation, the process of paying attention to the means of performing certain actions and trying to imitate them. Thus, it is not affordance perception that gives us access to modal knowledge, it is intrapersonal imitation. Consequently, a modal epistemology based on McClelland's account collapses into an epistemology of ability based on intrapersonal imitation.

4.5 Conclusion

In this chapter I have presented and analyzed different accounts of affordance perception and different ways of elaborating an epistemology of possibility based on them. After that I argued that none of those accounts work as a base for an epistemology of possibility. Some accounts, namely those that see action and perception as closely linked, require ability knowledge for conscious perception. This is problematic, because in doing so those accounts run into the problem of circularity when trying to explain possibility knowledge. If we perceive because of ability knowledge and if that perception explains our possibility knowledge, then we have explained modal knowledge in terms of other modal knowledge. Accounts of affordance perception that do not make this requirement fare better, in that they do not run into circularity issues, but they fail because affordance perception by itself cannot provide possibility knowledge. This is because it gets things in the wrong order, since we usually *see* a tree as climbable only when we already have the knowledge that we can climb trees in general, it is object-dependent whereas possibility knowledge is not, and it cannot explain how we keep our possibility knowledge in the absence of perceiving. A potential way of escaping these problems comes from McClelland (2019), whose account sees affordances as something that calls us into action by activating our motor systems. Thus, affordances can make us go into exploratory behaviors and learn about our abilities. However, I have argued that in order to learn about our abilities, we need to form a special type of representation and that we do not do this through affordance perception, but rather through *intrapersonal imitation*.

Chapter 5

Ability Knowledge

5.1 Introduction

In the preceding chapters I have analyzed and critiqued three types of perception-based modal epistemologies. I have argued that they all fall prey to the problem of modal epistemic friction (*PMEF*) and that they are unable to explain the cases of possibility knowledge they are interested in. Since Strohminger rejects the claim that our perceptual episodes have modal content, her account offers us nothing in the way of explaining how exactly it is that we go from perceiving modality to modal knowledge, which is far from a trivial jump. Thus she must either opt for a theory of perception where modality is part of the content of our perceptual episodes (sensorimotor approach or affordance theory) or abandon the idea that perception can lead us to modal knowledge. Nanay's account, on the other hand, holds that some of our perceptual states include modal content. For Nanay, our perception is intimately linked with action and in addition to our usual perceptual representations there is something he calls 'pragmatic representations', i.e. representations for action, which encode all the information that is necessary for the successful performance of the desired action. Some of that information is modal in nature. However, Nanay's account proves to be vulnerable to the *PMEF* as well. First of all, there is a difference between perceiving modalities and having modal knowledge and it is not at all clear how perception itself can guide us from one to the other. Second, Nanay rejects the idea that our perceptual states track modal properties out there in the world and he only accepts the claim that our perceptual states *represent* modal properties. But then what is it that makes successful cases of property attribution successful? This is a version of *PMEF* and, again, it seems like this modal epistemology would have to be changed from a perception-based one into another type of modal epistemology. Finally, for affordance accounts, we have two types of theories: the sensorimotor approaches and the classic affordance theories. They run into some issues as well. If we take sensorimotor theories, there seems to be a circularity problem going on, since they take it that percep-

tion requires some form of ability knowledge. Since ability knowledge is already a form of modal knowledge, they end up explaining modal knowledge in terms of other modal knowledge. In the case of the affordance theories it is not clear how we can explain the fact that we retain the modal knowledge of my ability to climb the tree in the absence of any perception of that tree. What is more, the modal knowledge is general, whereas the perception seems to be subject relative. Thus, it seems like perception-based modal epistemologies cannot explain the cases they set out to explain.

In this chapter I will argue that those cases are best explained by an ability-based approach to modal knowledge. We know that it is possible for us to climb the tree not because we perceive the tree as climbable, but because we have knowledge of our ability to climb trees. In the first section I will say a few words about abilities and how one should think about them. This will be an introductory and expository part of the chapter. In the second section I will present Barbara Vetter’s promising account of ability knowledge and I will argue that it is insufficient. In light of this, in the final section I will add to this account by arguing that our knowledge of abilities is formed in and through bodily awareness.

5.2 Abilities

What are abilities? Whenever we talk we tend to ascribe abilities to other people, we tend to say that ‘X can do such and such’, ‘Y is able to such and such’, etc. (Maier, 2022). But what are abilities? Perhaps the best way to figure out what we are talking about with abilities is to compare them to dispositions (Maier, 2022). Dispositions are properties that refer to certain tendencies of objects, like ‘being fragile’, being disposed to break when struck’ etc. Like abilities, they exist even when not manifest: a thin branch has the disposition to bend even it is not presently bent. They are, like abilities, connected to agency: ‘to the susceptibility of things to be acted on in certain ways’ (Maier, 2022), the actions we ascribe to things in the world (the sun rises), and the powers of agents. All of these aspects relate to powers, either of agents or of objects, so dispositions can be said to be a specific kind of *power* (Maier, 2022).

As Maier notes, abilities are also a kind of power, but we have to consider them as belonging to a different subset of the class of powers, for a number of reasons. First, unlike dispositions, abilities refer only to the powers of *agents* (Maier, 2022). Things cannot have abilities, only agents can, so abilities are the powers of agents (Maier, 2022). But this is not a sufficient condition for calling something an ability, because people have powers that are not abilities, like the disposition to blush (Maier, 2022). So we need a further condition: abilities relate agents to actions, they are powers through which agents perform actions (Maier, 2022). However, we should be careful to specify that abilities are powers through which agents perform actions that they initiate and that are performed in a controlled manner. In other words, besides being powers of agents and relating agents to actions, they relate agents to actions that they

initiate and that they have control over.

Maier argues that can distinguish between two types of ability: specific and general. Specific abilities refer to scenarios in which the agent is in a position to manifest their ability, while general abilities refer to scenarios in which an agent has an ability that they cannot showcase (Maier, 2022). For instance, take the example of two tennis players, one who is on the tennis court, with a racquet in hand, in a position to serve, and the other in the city center, eating at a restaurant. While both have the general ability to serve a tennis ball, only the first one has the specific ability to do it (Maier, 2022).

Another important thing that has to be mentioned is that when we deal with abilities we are not dealing with simple possibilities, but with a more restricted form of modality, which we will call 'agentive modality'. This is because abilities do not conform to the logic of the possibility ' $\Diamond$ ', and more specifically do not validate axiom T: $p \rightarrow \Diamond p$ (Vetter, 2024). In the modal logic of possibility, if something is actually true, then it is possibly true: if I score a three pointer on a basketball court, then it's (simply) possible that I score a three pointer on a basketball court. However, abilities do not work like that, since a case in which I perform an action successfully is not necessarily a case in which I have the ability to perform that action. In other words, the mere fact that I do something successfully one time does not imply that I have the ability to do that thing. This is because an ability requires a certain level of control and the possession of some kind of know-how, which makes it more likely in general that I will perform the action successfully and does not rely on luck. Of course, there is a sense in which the concept of ability can be used in a way that is compatible with simple possibility, where 'I am able to speak Japanese' indicates that it's possible, in principle, that I learn Japanese and get to speak it. It is open for me to speak Japanese. This type of ability is usually called *simple ability*. But when we speak of abilities we usually mean it in a stronger sense, where the agent initiates the action and has a sufficient degree of control over it. Those abilities are usually called 'robust abilities' and it is robust abilities we are interested in when we are analyzing abilities and ability knowledge. It seems like we have a great deal of knowledge about our robust, and this is not surprising. After all, we need to make choices on how to act all the time and in order to do this we need to know what our robust, general abilities are. It's not surprising at all that we have developed the capacity to know what our abilities are, since we can't survive without making choices and decisions on how to act. Specific and simple abilities are not enough for this task (Vetter, 2024). First, simple abilities tell us nothing about what it is likely for us to achieve, since it gives us no information about the probability of successfully performing an action (Vetter, 2024). I know that, in principle, I can speak Japanese, but that does not tell me anything about how likely it is that I will be able to do it successfully. Second, specific abilities seem to be to fleeting and context-dependent for them to be useful in decision-making (Vetter, 2024)¹. What we need is knowledge of

¹There are philosophers of ability who disagree with the claim that general abilities are essential for decision-making, and argue instead that specific abilities are more important

robust and general abilities (Vetter, 2024).

In the next section I will present and critically assess Barbara Vetter's agency-based account of abilities and, relatedly, her ability-based modal epistemology. I will argue that Vetter is on the right track, but that her account would benefit from an addition that I will provide. In short, the appeal to agency doesn't really doesn't explain our ability knowledge completely and it is still too vague and general. To see how we gain ability in more detail, we need to look towards the body and our relationship of awareness to it.

5.3 An Agency-based Epistemology

For Vetter, modal epistemology has so far focused mainly on types of modality that conform to the logic of possibility and necessity. Even though there is a lot of interest in ordinary, day to day modalities nowadays, theorists still see those modal species as restrictions on the abstract modal space that is studied in philosophy and logic (Vetter, 2024). Thus, knowledge of ordinary modality is seen as knowledge of restricted possibilities and necessities, to be analyzed through the same metaphysics, logic and epistemology (Vetter, 2024). However, for Barbara Vetter, there is a species of modality that does not lend itself to analysis with boxes and diamonds, namely agentive modality, concerned with ability. This type of modality has a different logic, as I have explained in the previous section, and in particular it doesn't validate axiom T, unlike simple possibility. I can, in principle, speak Japanese, I can utter a sentence in Japanese, but that does not entail that I have the ability to speak this language. In order to have such an ability, I need to be able to exhibit some level of control over the performance of that action (Vetter, 2024). A successful utterance of a Japanese sentence might just be a case in which I simply repeated a sentence I heard from a Japanese person or from a person who knows how to speak the language. Thus, one successful performance does not entail the possession of an ability to perform the action in question (Vetter, 2024). Therefore, there is a peculiar species of modality which does not conform to the abstract models we normally use in philosophy and logic (Vetter, 2024). Corresponding to it, then, is also a species of modal knowledge which needs special treatment (Vetter, 2024). As Vetter says, it would be a pretty big lacuna if our modal epistemology couldn't account for an entire subset of modal knowledge. So we need an epistemology that can explain how we come to know our abilities and, as Vetter argues, none of the available options seem to be up to the task. In the following subsection I will explain why Vetter thinks perception, counterfactual and similarity theories cannot account for our ability knowledge.

(Maier, 2015; Spencer, 2017; Schwarz, 2020). However, even if this is true, it doesn't affect my main argument.

5.3.1 Why Other Theories Fail

2

Perception-based theories

As we have seen in the previous chapters, perception-based theories claim that perception can give us epistemic aspect to the modal space, i.e. they can tell us something about what is possible and necessary. This usually goes well with theories of perception in which it reveals possibilities for action and where perception is intimately linked with action, like sensorimotor or affordance approaches. According to those accounts, we see trees as climbable, mugs as reachable, chairs as sit-on-able, etc. So there is a close link between perception and our agency in the world and affordances seem to be intimately connected to our abilities (Vetter, 2024). After all, they reveal the tree as climbable for the agent only if the agent has the ability to climb it. So it might be argued that by revealing affordances, perception also reveals our abilities.

However, as Vetter notes, the problem lies in the fact that affordances reveal object-dependent abilities, whereas our account should account for general and robust abilities (Vetter, 2024). Object-dependent abilities are those that are extrinsic to the subject performing the action, in that they can be lost and gained with changes in the object upon which the action is performed (Vetter, 2024). I lose the ability to climb that tree if it burns down, and my possession of that ability seems to depend on the integrity of the tree in front of me. As Vetter notes, affordances indeed reveal such abilities: when I look at the tree I become aware of my ability to climb this particular tree. But it is hard to see how affordance perception could teach us about our robust and general abilities, like my ability to climb trees *in general* (Vetter, 2024). Defenders of this account could reply by supplementing perception with some form of implicit reasoning (Vetter, 2024). This way, we generalize from our knowledge of object-dependent abilities, revealed to us through affordance perception, to knowledge of full-blown robust and general abilities. I see that I can climb this tree, that tree, and so on, and at some point I make a generalization and form the belief that I can climb trees in general (Vetter, 2024). As Vetter points out, the first problem would be that this solution accounts only for a subset of our abilities, namely those that are tied to external and perceivable objects. But our abilities' range is bigger. I have the ability to sing, to count, to speak, etc. Such abilities cannot be explained by affordances, since they seem to only be able to reveal what can be done to and with external objects that we can perceive (Vetter, 2024). The second issue is that the order of things is different (Vetter, 2024). When I see that I can climb the tree this is at least partially because I have some prior knowledge of my abilities, my ability to climb trees, or to climb in general, or to move my body in certain ways (Vetter, 2024). While affordances might give us new information about some aspects of our abilities, affordance

²The information conveyed in this chapter is not mine, it is taken from Vetter (2024).

perception cannot form the basis of our ability knowledge, since it itself will be mediated by prior ability knowledge (Vetter, 2024).

Conditional-based theories

As Vetter notes, conditional-based theories link abilities to counterfactuals: I am able to ϕ only if it is true that I would ϕ if I tried to ϕ . Consequently, counterfactual knowledge could provide knowledge of ability (Vetter, 2024). This is closely linked to the counterfactual-based accounts in modal epistemology, which claim that the mechanisms used for gaining modal knowledge are the same as those we use to evaluate counterfactuals (Vetter, 2024). More specifically, the proponents of such theories talk about our ability to formulate and develop off-line hypotheses through imagination (Vetter, 2024). Williamson provides an example that might be used to illustrate the connection between counterfactual imagination and ability knowledge (Vetter, 2024). Imagine a hunter who tries to evaluate whether he has the ability to jump over the stream that lies ahead of him (Vetter, 2024). He imagines trying to jump the river and develops this counterfactual scenario in his imagination (Vetter, 2024). If it develops into him successfully getting to the other side then he will judge himself able to perform this task, whereas if it doesn't then he will try something else, as he will not be confident that he possesses such an ability (Vetter, 2024).

One problem with this account is that it is not powerful enough to account for our abilities (Vetter, 2024). From the proposition 'If I tried ϕ , I would ϕ ' it does not follow that I am able to ϕ , unless I have the ability to try ϕ -ing (Vetter, 2024). It might be true that 'if I tried to climb a mountain, I would climb the mountain', but it doesn't follow that I have the robust ability of climbing the mountain (Vetter, 2024). As we have seen, I might have the simple ability to do that, but a robust ability necessitates a greater level of control (Vetter, 2024). Therefore, as Vetter argues, simply knowing that 'If I tried ϕ , I would ϕ ' does not by itself give me the knowledge that I am able to ϕ . Thus, knowledge of conditionals is not sufficient by itself and requires prior knowledge of abilities in order to deliver new ability knowledge (Vetter, 2024). Consequently, it cannot be that counterfactual knowledge is the source of ability knowledge, even though it sometimes plays a non-generative role in acquiring that kind of knowledge (Vetter, 2024).

Another problem with the account based on counterfactual knowledge is that it seems to come in conflict with some data about cognitive development (Vetter, 2024). Whereas ability knowledge is cognitively easy and comes earlier in development, the ability to evaluate counterfactual scenarios comes later and it is harder (Vetter, 2024). Studies show that adult-level counterfactual reasoning ability is not present in children under the age of eight, and even in cases in which younger children use counterfactuals, they do it differently (Rafetseder, Cristi-Vargas and Perner 2010; Rafetseder, Schwitalla and Perner 2013; Perner and Rafetseder 2011; Leahy, Rafetseder and Perner 2014; Nyhout, Henke and Ganea 2019). In contrast, ability-ascription is an something that children master early on and they seem to do very well at it (Papafragou, 1998; Cournane,

2020). Younger children are certainly capable of using ability information in goal-directed activities, when they have to choose means to an end and decide what to do in order to reach a certain goal. Counterfactual reasoning is more difficult to master and they do it later on (Vetter, 2024). Therefore, it isn't plausible to say that we gain ability knowledge through our ability to evaluate counterfactuals (Vetter, 2024).

Similarity theories

According to similarity theorists, we come to know unactualized possibilities through similarity judgments. Let's say I see a chair breaking, in actuality. First, I use axiom T to infer possibility from actuality and judge that it is possible for this chair to break. Then I remember that I have an object at home that is the same as this chair in the relevant respects and on the basis of this I reason that since it was possible for this chair to break it is possible for my unbroken chair at home to break as well, because they are relevantly similar.

First of all, the account is ill-suited for ability knowledge from the start, since we don't gain abilities or ability knowledge by seeing the actions of others (Vetter, 2024). As Vetter notes, at most we could say that we learn about our abilities on the basis of exercising them, and there is some kind of similarity reasoning involved in that. But even so, the focus of similarity theories is on alethic modality, as evidenced by the fact that axiom T plays a crucial role in this account (Vetter, 2024). However, as Vetter points out, agential modality is not alethic and it does not validate axiom T. As we have seen, a one-off success does not entail an ability to successfully perform something, since that latter thing requires a greater level of control. The similarity theorist might reply that axiom T can play a preliminary role (Vetter, 2024). We can think of simple abilities as simple possibilities and of robust abilities as possibilities under our control, so then we might say that in order to learn about our robust abilities, we need to learn about our simple abilities, so we need to learn simple possibilities, which validate axiom T (Vetter, 2024).

However, it is highly implausible that we come to know simple possibilities before we come to know our own abilities (Vetter, 2024). First of all, knowing about what we are able to do is necessary for successful action and successful action is crucial for our survival, so it must be that learning about our abilities is a process that is easy to master and that is developed earlier in life (Vetter, 2024). The same cannot be said of simple possibility, Vetter argues. It is not as crucial for our survival to know that we have the simple ability to speak Finnish (Vetter, 2024). If it isn't a robust ability, it does not give us useful information for planning or choosing between actions (Vetter, 2024). Simple possibility comes in handy when we need to be prepared for any contingency or when we need to solve a complicated problem, but those types of abilities are developed later in life and are more recent, evolutionarily speaking (Vetter, 2024). Furthermore, if we think about our everyday lives, we rarely encounter claims of simple possibility, but our lives abound in ability claims (Vetter, 2024).

5.3.2 Vetter's account

Vetter draws a number of lessons from those critiques, and we can think of them as desiderata for any epistemology of ability. First of all, an epistemology of ability should deal with the 'control' aspect of abilities, as we have seen counterfactual-based theories and similarity theories don't. Second, the account of ability knowledge should focus on the agent herself, not on the objects on which she performs her actions, like perception theories do. Finally, an epistemology of ability should not appeal to modal logic, counterfactuals or other abstract capacities/mechanisms, but to processes and capacities that are easy and developmentally early. With this in mind, Vetter proposes an agency-based epistemology of ability, whereby agency is at the root of our acquisition of ability knowledge (Vetter, 2024).

We engage in actions all the time, from a very young age. At first glance, it is intuitive to say that we try to do something, we repeat that thing and then, after a sufficient number of repetitions, we get it right. And after getting it right a sufficient number of times, we form the belief that we are able to do that thing. This is a naive flat-footed empiricist epistemology of ability and it says that our learning about abilities is a matter of repeated success. However, success cannot be all there is to the process of acquiring ability knowledge (Vetter, 2024). After all, I blush repeatedly and with success, and I might even win the lottery regularly, but those things are not what we would call robust abilities. So success must be supplemented with something else. Vetter suggests that the action must be an agentive phenomenon. Abilities are different from disposition or a streak of luck in that they exhibit agency: when I attribute an ability to myself I see the successful action as my own doing, as initiated and controlled or led to success by me (Vetter, 2023). These feelings, the feeling that the action was initiated at my will, that it was controlled and led to success by me and that it required some degree of effort (the greater the skill, the lesser the effort, and vice versa), make up what is called agentive phenomenology. Thus, we gain ability knowledge after observing that we have successfully performed an action with the feeling of agency a sufficient amount of times. Agentive phenomenology needs to be supplemented with the repeated success condition, because without it it is non-discriminating: it might make me feel like I am the agent of the action without isolating which aspects I am in control of. The repeated success tells me which aspects of the action I can reliably control. Furthermore, my agentive phenomenology by itself may be unreliable at times, a one-off success might be accompanied by a great feeling of agency and deemed a sign of ability. On the other hand, repeated success needs to be supplemented by agentive phenomenology, since without it we cannot distinguish between actions that are mere happenings and actions that we initiate and control, so we would consider mere dispositions or situations of luck or coincidences abilities.

Vetter's account presents *prima facie* plausibility and it manages to fulfill all three desiderata. It focuses on the agent herself, it addresses the control condition and it explains ability knowledge by appeal to capacities or mechanisms that are easy to get and developmentally early.

One problem with Vetter's account is that it falls prey to the *PMEF*. This is because we can ask the following question: How is it that agentive phenomenology, coupled with repeated success, lead us to ability knowledge? Vetter might reply that we can get out of this conundrum by adopting an reliabilist reading of agentive phenomenology. On this reading, the sense of agency is a mechanism that, among other things, reliably tracks our agency. Since part of our agency is taken up by the exercise of our abilities, agentive phenomenology reliably tracks our abilities. Since it is fallible, it must be constrained by something, so that it doesn't systematically misfire. That something is the repeated success condition, that prevents scenarios in which our feeling of agency wrongly attributes abilities to the agent: for instance, the repeated success condition prevents cases in which we attribute an ability to ourselves on the basis of a one-off success. It's very plausible that those elements are present in the process of ability knowledge acquisition, but one might question whether they are the essential components of this process. When the agent knows they have an ability, for instance playing the violin, they know that they can initiate the activity of playing a violin, *any violin*, in *any* environment where such a musical instrument is present. Therefore, when the agent knows they have an ability, they know that they can initiate any action-token of the action-type corresponding to that ability. This is different from the situation in which the agent represents their own actions without knowing about their abilities. When the agent doesn't know what they can do, they will represent their own actions as particulars that don't belong to any action type or as a stream of behaviour. Furthermore, we don't have a notion of *being able* to do something irrespective of the environment that triggered that action. But then a question remains: how are we led from not knowing about our abilities, from the stream of behaviour, to knowing that we can do a type of action without actually being engaged in performing a token of that action-type?

In order to know that we can do an action of the type ϕ in the absence of a scenario in which we actually do it, we need to have a representation of the means to do ϕ , a representation that is to some degree disconnected from the environment that triggered the first performance of an instance of the action-type ϕ . Such a representation exhibits two crucial properties for ability knowledge: information about how to perform the action, and a representation of the action type ϕ , as opposed to information about a specific instance of that action type. The first element is essential because in order to meet the control demand of abilities we need to have an idea about the technique required for the performance of the actions that correspond to that ability. The feeling of agency itself cannot arise without this prior condition, because we cannot intentionally initiate something that we have no idea how to do and because there is no feeling of control without the awareness that we have some grasp on the means of achieving the performance of an action. The second element is crucial because we are interested in knowledge of robust and general abilities. Those are abilities that refer to the successful and controlled performance of action types. When I am talking about my ability to write, I am talking about my ability to perform, in a successful and controlled fashion, most of the particular instances of the

action type of writing. This means that I am able to write independently of most variations in my environment, i.e. my ability to write is not limited to the environments in which I learned the ability. Thus, knowing I am able to perform an action type ϕ requires that I have a representation of the means to do most actions of type ϕ in all the environments that afford the possibility of writing, i.e. a representation where the performance of ϕ is decoupled from the environment that triggered the first performance of an action of type ϕ . In the first phases of learning an ability our representation of the means to perform actions related to that ability lacks generality and is most likely heavily dependent on context. In order to attain the ability and gain knowledge of it we must update this representation through learning, such that it becomes general and as context-independent as possible (of course, abilities cannot be fully context-independent, since they need environments that afford actions of specific types, i.e. I cannot play basketball in outer space). Sometimes we even measure the degree of ability by the presence of ability in as many contexts as possible. We consider a football player who plays well every week and in every weather condition better and more able than a football player who has a good game once every two months and who tends to play worse when it rains. This example does not have to do directly with the ability knowledge of the two players, but it does point to the fact that we tend to model ability in this fashion and that it plays an important role in attributing abilities to ourself and, thus, in the acquisition of ability knowledge.

This is not to say that the feeling of control and initiation and the repeated success condition do not play any role in ability knowledge acquisition, it is just to say that the heavy lifting in ability knowledge acquisition is done by the process whereby we come to form the type of representation described above. And it is hard to see how agential phenomenology and repeated success can fulfill this crucial role in the process of gaining ability knowledge. But if they can't, then what explains this phenomenon? An interesting suggestion comes from the work of Ellen Fridland on imitation. She identifies *imitation* as a locus of significance for discussing the difference between human and non-human animals (Fridland, 2013). While the act of imitation has been observed in both human and non-human animals, human animals tend to have a different relation to it. What Fridland contends is that humans exhibit a 'means-centric orientation', as opposed to a merely instrumental orientation, towards the imitation of a certain behaviour. In other words, while non-human animals have an almost exclusive interest in the goal of the imitated behavior, humans tend to focus more on the *means* of getting to that goal, that is we tend to focus on and attend to the ways in which the imitated action is performed (Fridland, 2013). This is what keeps us especially inclined to imitate more than other animals and, as Fridland says, it is one of the ways in which we acquire certain important 'practical and cultural competencies'. This usually comes from *interpersonal imitation*, i.e. the imitation of others, whereby we observe the actions of other people and try to replicate it. When doing so, we are interested not only and not primarily in the goal of the replicated action, but also in the technique by which the replicated action is performed. However, Fridland notes that there is a point at which we

start engaging in *intrapersonal imitation*, whereby we apply this means-oriented imitation activity to our own actions. This is the moment when a transition in cognition and learning occurs, whereby the child turns their focus and attention to the particular and detailed way in which they perform their own actions (Fridland, 2013). It is only by appeal to this turn that we can explain how the human child gets to the point where they can pay attention to, control, guide and fine-tune their actions. And, more importantly it is only by going through this transition that the child develops their sense of agency. In Fridland’s words:

By embracing the shift from the interpersonal means-centric orientation to the intrapersonal means-centric orientation, we find ourselves in a position to explain how it is that a child first begins to control, guide, attend to, and refine her own actions. By embracing this transition, we are in a position to explain how a child’s own abilities and behaviors become a “problem space” for her. And once we have done this, as I will argue below, we are in a position to explain the birth of the agentive features characteristic of cognition. (Fridland, 2013, p. 216)

It is at this point that one’s actions become ends in themselves and we start to attend to the detailed manner in which we perform our own actions. This has several consequences: it allows us to take control of and guide our own actions, it allows us to improve our skills and it allows us to better differentiate between different action components. Furthermore, as Fridland points out, it allows us to engage in trial-and-error, which is hugely important for the development of abilities and ability knowledge. This is because it allows us to explore our own actions by exploring the ways in which they can be performed. By engaging in trial-and-error, we get to see all the variations that can be exhibited by our actions. In this way, we can start to group particular action instances into action types and decouple action types from the initial trigger environment and from any specific or particular environment. In other words, applying the means-centric orientation to our own actions leads to the formation of a representation of what we can do in general and robustly. Thus, intrapersonal imitation can explain the formation of the type of representation that I argued above is essential for the acquisition of ability knowledge. But more so than that it underpins our sense of agency. This is because, as we have seen, this turn towards attending to the techniques involved in the performance of our own actions gives rise to the experience of controlling and guiding our actions and, more generally, to the experience of seeing actions as our own, which agency is all about. Thus, it seems that intra-personal imitation is the phenomenon that guides us towards ability knowledge by allowing us to form the required type of representation. Agentive phenomenology and repeated success, which is a consequence of the trial and error process, are brought about by intra-personal imitation, so it seems like *it* is the one doing the heavy lifting towards the acquisition of ability knowledge. As such, the sense of agency becomes a surface element in the process of ability knowledge acquisition where the heavy lifting is done by the phenomenon of

intrapersonal imitation. Vetter’s account appears to fall prey to the *PMEF* and collapse into an epistemology of ability that is based on intrapersonal imitation. This is mainly because intrapersonal imitation is more basic than the sense of agency or the repeated success condition. In order to account for it, we need to appeal to something even more fundamental and basic. In the next section I turn my attention towards bodily awareness.

5.4 Bodily Awareness

I know that the action is mine because I made or didn’t make my body do it, in a certain sense. Even in patients who are paralyzed and lose the ability to perform actions, a sense of ability is maintained, albeit in a negative sense (de Vignemont, 2018). They still feel like their body is theirs and they maintain a sense of what they can or can’t do (mostly of what they can’t do). If the sense of ability is so intimately connected to this sense of agency then paralyzed patients should also lose their sense of agency, but they don’t. Thus, we should perhaps look to something different to explain our relationship to ability knowledge.

In the previous section I argued that Vetter’s account falls prey to the *PMEF* because there is another phenomenon, namely intrapersonal imitation, which does the work of guiding us to ability knowledge, and it cannot be accounted for in terms of the sense of agency. In fact, the sense of agency itself arises out of intrapersonal imitation. Thus, the ability knowledge theorist has two options: they either opt for an account based on intrapersonal imitation, or they look to something more basic, that can account for intrapersonal imitation. The first option won’t do. The reason for this is that an ability epistemology still has to account for the feeling that the action corresponding to the ability in question is *mine*. This is crucial to ability knowledge because, as I have noted before, the feeling of mine-ness makes the difference between how we relate to abilities and how we relate to dispositions or instances of luck. Furthermore, we cannot even start intrapersonal imitation without experiencing the bodily movements that make up an action as *ours*. However, the feeling of agency comes out of intrapersonal imitation, so we cannot appeal to the feeling of agency to explain this phenomenon. We must look for something more basic than the sense of agency, something that can both account for intrapersonal imitation and our feeling of mine-ness.

One promising suggestion comes, as I alluded to above, from the philosophy of bodily awareness, and more specifically from the work of Frederique de Vignemont. Our body is a peculiar object, since on the one hand we seem to always be in contact with it, but on the other hand we know surprisingly little about it (de Vignemont, 2024). Thus, it’s the object we know and are connected to best, but at the same time it can prove to be the most elusive object we are in contact with (de Vignemont, 2024). De Vignemont is primarily interested in elucidating the nature of this peculiar relation we have with our bodies, but I think that her insights might help in elaborating a cognitively plausible, empiricist and ability-based modal epistemology.

For de Vignemont, we are in constant contact with our body and we elaborate representations of it all the time (de Vignemont, 2018; de Vignemont, 2024). Most theorists differentiate between two types of representations: body schemata and body image (de Vignemont, 2018). The first is split into two schemas: one that represents every movement and one that maps locations in the body with sensations that are felt at those locations (de Vignemont, 2024). However, as de Vignemont shows, this distinction doesn't really manage to account for various bodily disorders and suggests that we should drop it altogether and search for a different type of representation (de Vignemont, 2018).

One place to start is to reconsider the relationship between bodily representations and action. The functioning of our motor system requires bodily information all the time, from the intention to move to the actual performance of the movement (de Vignemont, 2018). As de Vignemont notes, the planning of action, for instance, requires information not only about what we should or will do, but also about the current state of our bodies and the control of action requires information about the location our bodies should be in and about the location our body is in. So the sensorimotor system builds a 'practical photograph of the body' throughout our lives (de Vignemont, 2018). This happens through the sensorimotor feedback that we receive when engaging in simple movements. For instance, infants kick their legs or perform repeated action on their bodies and explore the correspondence between visual information and proprioceptive information, i.e. information about the position of their body parts (de Vignemont, 2024). The sensory feedback allows the subject to adjust their bodily representations and as we grow we become more and more aware of some of those representations and, in normal situations, of our bodies (de Vignemont, 2018). The content of these representational states can be cashed out in different ways. One suggestion is to cash them out in terms of affordances, which we have talked about in the last chapter, and more specifically in terms of bodily affordances (de Vignemont, 2018). Bodily affordances 'tell' us, or rather show us, how we can use certain body parts and more precisely how we can move them (de Vignemont, 2018). There are two types of bodily affordances: postural affordances, which give us information about the position of body parts and how we can move them here and now (they form at time t and disappear at time $t+1$), and structural affordances, which tell us about longlasting properties of our bodies, like the size of my leg, which hasn't changed in a few years and which will not likely change in the future (de Vignemont, 2018).

Instead, de Vignemont proposes that we use a concept that she adopts from Millikan, namely *pushmi-pullyu representations* (PPR) (Millikan, 1995) to account for bodily representations. It is assumed that, depending on their relation to the world, representations can be of two types: descriptive, in which case they have truth conditions and they are said to have a mind-to-world direction of fit (because the representation must be changed in order to conform to the world), and directive, in which case they have success conditions and are said to have a world-to-mind direction of fit (because the world must be changed in order to correspond to the content of the representation). PPRs have both directions of fit, they direct action and at the same time they tell us something about the

world (Millikan, 1995). One example of this is the act of the priest declaring two people husband and wife. This act is both descriptive, because it describes the nature of the relationship between the two people, and directive, because it says that the priest is making it happen (de Vignemont, 2018). Representations of bodily affordances themselves are examples of PPRs, because they simultaneously tell you what you can do and what you should do. The motor system involves two kinds of models: the inverse model, which is tasked with representing the motor command that is necessary to achieve a desired outcome, given the bodily affordances of the agent, and the forward model, which has the role of predicting the outcome of the action in terms of the body performing it and to provide 'anticipatory control of movements' (Wolpert et. al, 2001). So the representations that guide bodily movements are bodily PPRs that describe the body in terms of what it can do (bodily affordances) and that 'direct bodily movement', i.e. that tell you what you should do (de Vignemont, 2018).

One key take-away from this account of bodily representations is that they compel us to act. Since we have them from infancy, they compel and direct us, unconsciously, to engage in certain behaviours the performance of which gives us sensory feedback. In other words, we are compelled to act and the consequences of the performance of those actions is that we learn about our bodies. If we find something interesting about an act, then we might feel compelled to repeat it. This is true of infants who are exploring the world for the first time and exhibit a lot of creativity, but also about adults, who often learn by doing things that have unintended consequences which hold important and interesting information. And this is where intrapersonal imitation comes in. Doing something that we find interesting and potentially important or that simply makes us curious is then bound to make us more attentive to the particularities of the act, and more specifically to its structure and the ways in which we can perform it. In other words, we start engaging seriously and more consciously with that action. In doing this, we become attentive to the particular ways in which that action can be brought about and *controlled*. And this is part of learning too, since we learn, intentionally or not, how to initiate and control that action. Finally, if we repeat that action successfully a sufficient number of times we build confidence that we can also do it in an unactualized scenario in the future, we believe it and we use those beliefs in our planning for the future.

But this is not enough, since we also need to feel like the actions are ours. This comes from what de Vignemont calls 'the sense of bodily ownership', the feeling that my body is mine, which is found in everyone, even in people who have lost the ability to move their body at will, for example people who are paralyzed (de Vignemont, 2018). This feeling is most probably evolutionarily old and it comes from the evolutionarily advantageous feeling that the body is the object on which the agents' survival depends most. This creates a certain level of *affective significance* that makes us identify, in a sense, with our own bodies or at the very least have the feeling that it is ours (de Vignemont, 2018).

Unlike Vetter's account, this framework allows us to explain intrapersonal imitation in terms of bodily awareness. Intrapersonal imitation only works because we are in intimate contact with our body and because we see it as *ours*.

Otherwise it's hard to see how we could apply the means-centric orientation to ourselves and our actions. To do this we need to see our actions as something performed by *our bodies*. It's not necessary to see our actions as ourselves or as our own doing and indeed we have seen that our sense of agency is something that develops out of intrapersonal imitation. But it is necessary to feel that what is doing the action is something that is ours, that we own. Furthermore, we pay attention to the means of performing an action by paying closer attention to the particular movements of our various body parts when engaging in that action. It is hard to see how we could do that without bodily awareness and the sense of bodily ownership. Without feeling that our body is ours we wouldn't be able to develop a sense of agency out of intrapersonal imitation, since we would relate to our body in the same way we relate to other people's bodies in *interpersonal imitation*. But this is precisely the locus of the difference between interpersonal imitation and intrapersonal imitation, namely the fact that the body we imitate is ours, it doesn't belong to someone else, that the body performing the action to be imitated is ours and that our body is engaging with the means of performing the action.

5.5 Conclusion

Thus, we gain knowledge about our abilities in the following way: we find actions that get our attention, make us attuned to the way they can be made to come about, we explore them and learn how they can be brought about, we repeat them and find that we have a certain level of success, and then we form the confidence and the belief that we can do those actions in the future when engaged in planning. This is made possible by our bodily awareness, by the fact that we constantly monitor and represent our body. However, those actions can lead to abilities only once we get to a point where we can have control over their performance and only in cases where we feel like our body performed those actions, as opposed to feeling like those actions were performed on our bodies. Like Vetter's account, mine also appeals to a part of our phenomenology, namely the sense of bodily ownership, but unlike hers, my account has the resources to deal with the phenomenon of intrapersonal imitation without falling prey to the *PMEF*.

Chapter 6

Conclusion

This thesis can be thought of as both a rebuttal of the idea that perception can be a genuine guide to modal knowledge and a defense of an approach to possibility knowledge that seeks to ground it in our ability knowledge. The argument for the latter is that it is an approach that explains possibility knowledge in a naturalistically feasible manner, by appealing to a capacity that is ordinary, easy and developmentally early. It is hard to find an earlier epistemic point of entry into the modal space than ability knowledge.

The dissertation has aimed to show two things. The first aim was to show the untenability of the idea that perception can be a guide to knowledge of possibility. The main argument for this is that whenever we try to elaborate a modal epistemology based on a theory of perception, our account falls prey to the *PMEF*. If we opt for a theory of perception where our perceptual representations do not have modal content, i.e. do not represent modal properties, like Strohminger does, then we can legitimately ask what it is that takes us from the perceptual episode to the piece of modal knowledge corresponding to it. This is because perception isn't directly modal, so we need an additional mechanism to take us from the non-modal perception to the instance of modal knowledge. Since the cases Strohminger is interested in are tied to what we are able to do, a plausible answer to this is that ability knowledge allows us to make that jump from perception to possibility knowledge. But then what guides us to the relevant justified true beliefs in the relevant possibilities is ability knowledge, not perception. If one opts for a theory of perception where perceptual episodes represent modal properties, then a modal epistemology based on it presents a variety of issues. First, if one opts for Nanay's action-based account of perception and build an epistemology of possibility on the basis of it, this account falls prey to the *PMEF*. This is because Nanay's approach rejects the claim that our perceptual mechanism reliably tracks modal properties that are out there in the world. But then we can ask what it is that makes our perception successful in the cases in which it gets things right (with respect to modal properties). Again, since his account of modal perception is tied to action, it seems like a plausible answer is ability knowledge. But then ability knowledge is the guide to modal

knowledge and this account collapses into an ability-based modal epistemology. If one opts for a sensorimotor approach to perception, then the modal epistemology based on it collapses into an ability-based epistemology of possibility. This is because on this approach perception itself requires prior ability knowledge. Finally, if one opts for a classic affordance perception approach then a modal epistemology of it is bound to not have the resources to explain possibility knowledge. This is because affordance perception is very object-relative, it provides information about what can be done to a specific object, whereas modal knowledge is general. If I have modal knowledge about a tree I have modal knowledge about trees in general. Furthermore, it cannot explain why we have modal knowledge in the absence of perception. An ability-based epistemology of possibility does not face these issues.

A second aim of this dissertation has been to defend an approach to possibility knowledge that is grounded in ability knowledge. In the final chapter I have compared two such approaches: Barbara Vetter's agency-based modal epistemology and my own approach, grounded in our bodily awareness and in the feeling of bodily ownership. I have argued that Vetter's account cannot explain a crucial step in the process of ability knowledge acquisition, namely the process whereby the agents forms a representation of what they can do, in the absence of the trigger environment and in a context independent manner. This is explained by the phenomenon of intrapersonal imitation, but Vetter's notion of the sense of agency cannot explain or account for this, since intrapersonal imitation is precisely that which gives rise to agential phenomenology. In contrast, my account is based on bodily awareness and the phenomenology of bodily ownership, which are more basic and earlier than intrapersonal imitation. The latter can be explained through by appeal to bodily awareness and the feeling of bodily ownership because we engage in personal imitation by engaging with the thing that we call 'our body'. We pay attention to the means of performing our actions by attending to the particular movements of our own body.

In sum, my thesis brings something new to the field of the epistemology of modality in two ways. First, some modal epistemologists are in the habit of claiming that perception is one of the ways in which we gain modal knowledge and just cite Strohming and Nanay without any critical analysis of their arguments. It seems like whereas before Strohming and Nanay modal epistemologists simply assumed that perception cannot deliver modal knowledge, whereas now they simply assume that it can. In my dissertation I show that the accounts of Strohming and Nanay fail and perception is *not* a modal epistemic guide. Second, I propose an original, ability-based epistemology of possibility, grounded in bodily awareness. This follows the recent trend of accounts that try to account for modal knowledge by appeal to ordinary, naturalistically/evolutionarily plausible capacities that are developmentally early and easy to attain, but it avoids some of the shortcomings of the other theories.

References

- Balcerak Jackson, M. (2018). Justification by imagination. In F. Macpherson & F. Dorsch (Eds.), *Perceptual imagination and perceptual memory* (pp. 209–226). Oxford University Press
- Byrne, A. (2007). Possibility and imagination. *Philosophical Perspectives*, 21, 125–144.
- Chalmers, D. (1995) “Facing up to the problem of consciousness”. *Journal of Consciousness Studies*, 2: 200–19.
- Chalmers, David J. (1996). *The Conscious Mind*. Oxford: Oxford University Press.
- Chemero, A. (2003). An Outline of a Theory of Affordances. *Ecological Psychology* 15:2, 181-195.
- Chemero, A. (2011). *Radical embodied cognitive Science*. MIT Press.
- Cournane, Ailís 2020: ‘Learning Modals: A Grammatical Perspective’. *Language and Linguistics Compass*, 14(10), pp. 1–22.
- Currie, G., & Ravenscroft, I. (2002). *Recreative minds*. Oxford University Press.
- de Vignemont, Frédérique (2018). *Mind the Body. An Exploration of Bodily Self-Awareness*. Oxford University Press.
- de Vignemont, Frédérique (2024). Bodily Awareness. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Stanford, CA.: The Metaphysics Research Lab, Stanford University, winter 2024 ed.
- Fiocco, M. (2007). Conceivability, imagination and modal knowledge. *Philosophy and Phenomenological Research*, 74, 364–380.
- Fischer, B., & Leon, F. (2017a). Introduction to Modal Epistemology After Rationalism. In B. Fischer, & F. Leon (Eds.) *Modal Epistemology After Rationalism*, (pp. 1-6). Cham: Springer.
- Fridland, E. (2013). Imitation, Skill Learning, and Conceptual Thought: An Embodied, Developmental Approach. In L. Swan (Ed.) *Origins of Mind*, (pp. 203–224). Springer.
- Gibson, James (2014). *The Ecological Approach to Visual Perception*. Psychology Press.
- Goldman, A.I. (1994). Naturalistic epistemology and reliabilism. *Midwest studies in philosophy*, 19, 301-320
- Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University Press.

- Gregory, D. (2020). Imagery and possibility. *Nous*, 54(4), 755–773.
- Grush, R. (2007). Skill theory v2.0: dispositions, emulation, and spatial perception. *Synthese*, 159, 389–416.
- Hale, Bob (2013). *Necessary Beings: An Essay on Ontology, Modality, and the Relations Between Them*. Oxford: Oxford University Press.
- Jenkins, C.S. (2013). Naturalistic Challenges to the A Priori. In A. Casullo, & J. C. Thurow (Eds.) *The A Priori in Philosophy*, (pp. 274-290). Oxford: Oxford University Press.
- Kung, P. (2014). You really do imagine it: Against error theories of imagination. *Nous*, 50, 90–120.
- Leahy, Brian, Eva Rafetseder, & Josef Perner 2014. ‘Basic Conditional Reasoning: How Children Mimic Counterfactual Reasoning’. *Studia Logica*, 102(4), pp. 793–810.
- Maier, John, 2015. The Agentive Modalities. *Philosophy and Phenomenological Research*, 90: 113–134.
- Maier, John (2022). Abilities. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Stanford, CA.: The Metaphysics Research Lab, Stanford University, fall 2022 ed.
- McClelland, T. (2019). Representing our options: The perception of affordances for bodily and mental action. *Journal of Consciousness Studies*, 26 (3-4), 155– 180.
- Millikan, Ruth Garrett (1995). Pushmi-pullyu Representations. *Philosophical Perspectives* 9:185-200
- Nanay, Bence (2010). Action-oriented Perception. *European Journal of Philosophy* 20(3)
- Nanay, Bence (2011a). Do We Sense Modalities with Our Sense Modalities?. *Ratio* 24(3): 299-310.
- Nanay, Bence (2011b). Do We See Apples as Edible? *Pacific Philosophical Quarterly* 92(3): 305 - 322.
- Nanay, Bence (2014). *Between Perception and Action*. Oxford University Press.
- Nichols, S., & Stich, S. P. (2003). *Mindreading: An integrated account of pretence, self-awareness, and understanding other minds*. Oxford University Press.
- Nolan, D. (2017). Naturalised Modal Epistemology. In B. Fischer, & F. Leon (Eds.) *Modal Epistemology After Rationalism*, vol. 378, (pp. 7-27). Cham: Springer.
- Nyhout, Angela, Lena Henke, and Patricia A. Ganea 2019: ‘Children’s Counterfactual Reasoning About Causally Overdetermined Events’. *Child Development*, 90(2), pp. 610–22.
- Perner, Josef, and Eva Rafetseder 2011: ‘Counterfactual and Other Forms of Conditional Reasoning: Children Lost in the Nearest Possible World’. In Christoph Hoerl, Teresa McCormack, and Sarah R. Beck (eds.), *Understanding Counterfactuals, Understanding Causation: Issues in Philosophy and Psychology*, pp. 90–109. Oxford: Oxford University Press.
- Priest, G. (2016). Thinking the impossible. *Philosophical Studies*.

- Rafetseder, Eva, Renate Cristi-Vargas, and Josef Perner 2010: ‘Counterfactual Reasoning: Developing a Sense of the “Nearest Possible World”’. *Child Development*, 81(1), pp. 376–89.
- Rafetseder, Eva, Maria Schwitalla, and Josef Perner 2013: ‘Counterfactual Reasoning: From Childhood to Adulthood’. *Journal of Experimental Child Psychology*, 114(3), pp. 389–404.
- Rysiew, P. (2020). Naturalism in epistemology. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Stanford, CA.: The Metaphysics Research Lab, Stanford University, spring 2020 ed.
- Salmon, Nathan (2005). *Reference and Essence*. 2nd ed. Amherst, NY: Prometheus Books.
- Schwarz, Wolfgang, 2020. Ability and Possibility. *Philosophers’ Imprint*, 20(6)
- Shapiro, L. (2010). *Embodied cognition*. Routledge.
- Siegel, S. (2006). Which Properties are Represented in Perception?. in T. Gendler and J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press, pp. 481–503, at p. 499.
- Spencer, Jack, 2017. Able to Do the Impossible. *Mind*, 126: 466–497.
- Strohming, Margot (2015). Perceptual Knowledge of Nonactual Possibilities. *Philosophical Perspectives* 29 (1): 363–375.
- Vaidya, A.J. (2017). Modal Knowledge: Beyond Rationalism and Empiricism. In B. Fischer, & F. Leon (Eds.) *Modal Epistemology After Rationalism*, vol. 378, (pp. 85–114). Cham: Springer.
- Vaidya, A. J., & Wallner, M. (2018). The epistemology of modality and the problem of modal epistemic friction. *Synthese*, 198(S8), 1909–1935. <https://doi.org/10.1007/s11229-018-1860-2>
- Vaidya, A. J., & Prelević, D. (2023). *Epistemology of modality and philosophical methodology*. Taylor & Francis.
- Vaidya, A. J., Mallozzi, A. & Wallner, M. (2023). The Epistemology of Modality. In E. N. Zalta (Ed.) *Stanford Encyclopedia of Philosophy*. Stanford, CA.: The Metaphysics Research Lab, Stanford University, winter 2016 ed.
- Vetter, B. (2015). *Potentiality: From Dispositions to Modality*. Oxford University Press.
- Vetter, B. (2023). An Agency-Based Epistemology of Modality. In A.J. Vaidya, & D. Prelević (Eds.) *Epistemology of modality and philosophical methodology*. Taylor & Francis.
- Vetter, B. (2024). Abilities and the Epistemology of Ordinary Modality . *Mind* 133 (532): 1001–1027
- Ward, D., Roberts, T., Clark, A. (2011). Knowing what we can do: actions, intentions, and the construction of phenomenal experience. *Synthese*, 181 (3), 375–394.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.
- Wolpert, D. M., Ghahramani, Z., and Flanagan, J. R. (2001). Perspectives and problems in motor learning. *Trends in Cognitive Science*, 5, 487–94.