

Presentations of Value:

Evaluative Outlooks and Practical Reason

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Michael Ebling

August 2024

Acknowledgments

I would like to thank all the many teachers and friends who made this work possible. I especially want to thank Jon Garthoff for advice and discussion at every stage of this project; Claire Dartez, Georgi Gardiner, Linh Mac, and Troy Pollidori for helpful comments, conversations, and critiques on various drafts of these chapters; and my dissertation committee—Jon Garthoff, Georgie Gardiner, Kristina Gehrman, and Garriy Schteynberg—for their support, suggestions, and willingness to help me tackle this daunting task. I benefited also from conversations with Jordan Baker about many topics relevant to this project. Finally, I am indebted to Lisa Hill for much encouragement and support along the way.

No doubt there are others whose important contributions I have overlooked. To all of you, thank you.

Abstract

In this dissertation, I argue for an evaluative outlook account of human practical reason by developing a viable representational psychology that vindicates the following key claims. First, some mental states are evaluative representations with ineliminably evaluative representational content. Second, any successful explanation of a rational action must appeal to evaluative representations. Third, many evaluative representations are products of subrational processes and capacities. Fourth, in humans evaluative representations function to be elements in an overall evaluative understanding. And fifth, evaluative representations by nature have motivational efficacy. In addition to these five foundational claims, I add two more speculative points. Some evaluative representations are partly constituted by their phenomenal quality. And the phenomenal quality of an evaluative representation (if it has phenomenal quality) partly explains its motivational efficacy.

In the first chapter, I state and discuss these key claims and the motivations of the evaluative outlook model. In the second chapter, I sketch an account of representational mind, from perception to propositional thought, and I consider how representations can have evaluative content. In the third chapter, I consider, and ultimately reject, the claim made by some philosophers that emotions are perceptions of evaluative properties. The fourth chapter focuses on capacities for propositional thought and understanding, ultimately developing a notion of intention as a conative state with conceptual content that functions to guide other conative states. In the fifth and final chapter, I develop the idea that life's meaningfulness is an end that functions in practical reasoning to organize other ends.

Table of Contents

Chapter 1 Evaluative Outlooks.....	1
I. Introduction: Ethics, Psychology, and Practical Thought.....	1
II. The Evaluative Outlook Model of Practical Reason.....	4
III. On the Nature of Value.....	13
IV. Dialectical Activity, Desire, and Phenomenal Consciousness.....	20
V. The Way Ahead.....	33
Chapter 2 Representation, Evaluation, and Agency.....	36
I. Representational Mind.....	38
II. Evaluative Representational Content.....	54
III. The Space Between Perception and Thought.....	59
IV. Desire and the Problem of Receptivity.....	73
Chapter 3 Emotions, Value, and Phenomenology.....	80
I. A Perceptual Theory of Emotion.....	81
II. An Emotional Theory of Value.....	86
III. The Nature and Epistemic Role of Emotions.....	93
IV. Presentation and Phenomenal Consciousness.....	101
V. Necessarily Phenomenal Modes of Presentation.....	107
Chapter 4 Thought, Understanding, and Intention.....	117
I. Thought.....	119
II. Basic Understanding.....	130
III. Metacognition, Critical Reason, and Theoretical Understanding.....	141
IV. Intention and Rational Motivation.....	147
V. The Outward-Looking Nature of Practical Reason.....	160
Chapter 5 Meaning and Choiceworthiness.....	163
I. Meaning and Dialectical Activity.....	165
II. Evaluative Outlooks and Fundamental Values.....	173
III. The Structure and Content of Practical Reason.....	185
IV. Conclusion.....	196
References.....	198
Vita.....	203

Chapter 1

Evaluative Outlooks

I. Introduction: Ethics, Psychology, and Practical Thought

Philosophers are divided about the purview of practical thought. One camp is skeptical of the idea that a person's ends—understood as the objects of their desires and projects—can be subject to rational scrutiny. According to this kind of view, the heart wants what it wants, so to speak, and no amount of reasoning can change that, or even generate justified critiques of it. On the other side of the divide are those who hold that there is a legitimate rational standpoint from which one can critique the objects of a person's desires. On this view, some desires and goals are better than others, and sound practical thinking can, at least sometimes, actually succeed in revising those desires and goals according to objective standards. Let us call the first family of views *subjectivist*, and the second family *objectivist*. Of course, there are many complexities and variations within these general categories, including views that straddle the divide in various ways; nonetheless, I find the distinction generally helpful.

Philosophers might be inclined toward subjectivism about practical thought for several reasons. One concern has to do with worries about evaluative facts. According to a certain kind of ontology, values have no place in objective reality, and would be unknowable even if they did exist, and so any critique or revision of what a person finds valuable cannot be based on objective considerations. Another related concern that inclines some toward subjectivism has to do with motivation. Many philosophers find it doubtful that objective representations of reality, like beliefs, could by themselves move a person to act. And so, the thinking goes, there must be

some subjective element to practical thought that provides the motivational force of practical conclusions. Moreover, whatever that subjective element is, it will not be subject to rational scrutiny because it is not a representation of reality that could be correct or incorrect.

Objectivists, on the other hand, hold that there are evaluative facts, and that there are evaluative representations that can and do move agents to act. For example, Talbot Brewer (2009) argues that some mental states are *apprehensions of value*, and that as such they have a special role to play in practical reasoning and rational action. I shall follow Brewer in calling views like this *evaluative outlook* views.¹ One appeal of this kind of approach is that it draws parallels between the practical and the theoretical aspects of thought, trading on the (putatively) better-understood nature of the latter. Another appeal is that evaluative outlook accounts generally do a good job explaining common ethical practices and pre-theoretical ethical attitudes. For these reasons, I think that the burden of argument is on the subjectivist. Nevertheless, much work remains to develop an adequate understanding of these matters.

I take seriously the suggestion that there are mental states specifically attuned to value. In doing so, I take an epistemological and psychological approach to what is ultimately a problem in ethical theory. At stake is the nature of practical thought, and I propose to develop an understanding of that by reflecting both on relatively primitive psychological capacities, like perception and emotion, and on sophisticated capacities, like thought and critical reason. This methodology is justified by the impressive degree to which these more primitive psychological capacities themselves are integral parts of human thought. Human rational thought develops out of capacities for perception, emotion, memory, and so forth. In this dissertation, I argue for an evaluative outlook account of practical reason by developing a viable psychology that explains

¹ See also Scanlon (1998) and Tenenbaum (2007). See Shapiro (2014) for a general critique of this kind of view.

and vindicates the key claims of that view. I think that practical reason constitutively functions to respond to what in fact matters, in the same way that theoretical reason constitutively functions to respond to what is in fact the case. Ultimately, I leverage this idea to develop insights into the nature of value and the structure of practical thought.

In this first chapter, I elaborate the key claims and motivations of the evaluative outlook model, and I clarify the central concepts that I will rely on throughout the dissertation. This chapter, then, is largely stage-setting for what is to come, but laying out the relevant concepts and distinctions is no small task, and the philosophical upshots of this mapping of the territory will be of interest in their own right.

I aim to articulate and make progress on constitutive questions about fundamental kinds within ethics and philosophical psychology. Before I discuss the central concepts I employ, it will be helpful to reflect on the general methodology at work here.

Tyler Burge (2010a) defines the key notions of *constitutive conditions* and *natures*:

A constitutive question concerns necessary (or sufficient, or necessary and sufficient) conditions under which something is what it is. Such a question concerns conditions under which something has the *nature* that it has. As I conceive them, natures are approximately essences. I want, however, to push to the background many of the traditional metaphysical questions about essence.

Natures are associated with fundamental, or relatively fundamental, kinds or properties that have the potential to figure in systematic explanations. What counts as a relatively fundamental explanatory kind must be determined in the rough and tumble of explanation....

A constitutive question asks for necessary or sufficient conditions for something's being what it is or having the nature that it has. The relevant conditions are, however, a subset of all such necessary and sufficient conditions. To be an answer to a constitutive question, the answer must help *explain* something's having the nature that it has. (57–58)

I follow Burge in pursuing constitutive questions. The mark of success is improved understanding via illuminating explanation. For the most part, I begin with pretheoretical concepts that pick out explanatorily important kinds. I try to clarify the natures of these kinds by specifying conditions under which something is the kind of thing that it is. These specifications are neither stipulations nor mere linguistic definitions. They are subject to revision in the course of both rational reflection and empirical investigation.

In many cases, the concepts and distinctions I outline here will operate in the background in the chapters to come. In other cases, I explicitly develop the distinctions in detail later on. I recognize that many of the points I put forth here are controversial, and I nonetheless rarely stop to offer detailed arguments or critique opposing perspectives. I beg the reader's patience. For the most part, I think the most persuasive argument for the views I present is the ultimate fruitfulness of the explanations they make possible.

II. The Evaluative Outlook Model of Practical Reason

Certain models of practical reason attribute to agents representations of value, and claim that these evaluative representations play an indispensable role in rational action. Moreover, these evaluative representations are typically understood to include relatively primitive kinds of representational state. For example, Brewer (2009) presents what he calls the "evaluative outlook conception of desire," and he explicitly rejects understanding desire in terms of belief (24–25). The notion of desire here is meant to capture the idea that some representational mental states are evaluative *appearances* that "play a role in practical reasoning akin to the role played by

perceptions in theoretical reasoning” (25).² My view follows suit, but I try to state it in a way that does not center on desire, not just because desire is a philosophically fraught notion, but also because I aim to focus on a more general psychological phenomenon: action-guiding evaluative representation. Thus, we can understand the view that I call the *evaluative outlook model of practical reason* as an account of practical reason that attempts to explain rational action partly in terms of evaluative appearances. I think it is appropriate to use the term “desire” to refer to some of these appearances, but my aim is not to theorize about desire per se. In fact, I think it is possible to lay out the main commitments of the evaluative outlook model of practical reason without any reference to desire, a task to which I now turn.

The evaluative outlook model, as I understand it, involves five general theses:

- 1) Some mental states are evaluative representations with ineliminably evaluative representational content.
- 2) Any successful explanation of a rational action must appeal to evaluative representations.
- 3) Many evaluative representations are products of subrational processes and capacities.
- 4) At least in humans, evaluative representations function to be elements in an overall evaluative understanding, or conception of the good.
- 5) Evaluative representations by nature have motivational efficacy.

² Brewer traces his view to the work of Tim Scanlon (1998) and Dennis Stampe (1987), and he contrasts it with what he calls the “propositionalist” account of practical reason, according to which no mental states have irreducibly evaluative content.

These five theses do not quite exhaust the basic evaluative outlook model. As we shall see, there are important connections between these claims that the model insists on. But rather than trying to list all the connections between them as separate claims to which the view is committed, I will elaborate the main five claims with an eye toward bringing out how they all hang together.

Some mental states are evaluative representations with ineliminably evaluative representational content. This first claim says that the evaluative element of an evaluative representation is part of the representation's content, not some separate non-representational mental state or valence. This claim insists that value can be represented, and thus value is something about which one can be correct or incorrect. *Representation* is the functional reference, and attribution of properties, to aspects of reality by an individual. The basic norm of representation is *veridicality* (accuracy). That is, a representation is veridical (correct) when it refers to an actual aspect of reality that in fact has the properties the representation attributes to it.³ So when an evaluative representation presents some object as having an evaluative property, the representation is veridical if the object really has that property. Thus, in order to make sense of representations with genuinely evaluative content, there must be evaluative properties that can be represented, at least approximately, by the representational state in question.

In the next chapter I provide a more detailed account of representation and what it means for a representation to have evaluative content. My argument for this first claim in the evaluative

³ Representation should be contrasted with mere information registration, including functional information registration. Reliable covariance conveys information, as in "smoke conveys information about fire." But this is not in itself functional, nor is it representation, as I use the term. Functional information registration, as when an amoeba senses light in a way that contributes to biological functioning, does not suffice for representation. Mere sensory information registration, as when an organism detects changes in its environment via proximal stimulation of its sensory organs, is not in itself representational. Of course, all perceptual representation is mediated by sensory information registration. But representations constitutively function to be *correct*, and representation is a fundamental *psychological* kind (Burge 2010a: 9, 27).

outlook model lies in the explanatory power of that account. For now, suffice it to say that mental representations are perspectives on an objective reality, and evaluative representational content involves an attribution of evaluative properties to objects, activities, or states of affairs. The most primitive kind of evaluative (or perhaps proto-evaluative) representational content is the application of functional attributives (such as food) at the level of perception, which function to be goals in perceptual-level conative states. More sophisticated evaluative representations include evaluative beliefs and explicit evaluative judgments. These more sophisticated kinds of representation are essential for understanding rationality, which brings us to the second key claim in the evaluative outlook model.

Any successful explanation of a rational action must appeal to evaluative representations.

The second claim takes a stand on what counts as an action, and in so doing provides some additional reason to accept the first claim. The point of this second claim is to specify the target of theorizing about practical reason. A theory of practical reason must explain *actions* as opposed to mere behaviors. To be a rational action, in the fullest sense, it is not enough for a behavior to be goal-directed; the choiceworthiness of that goal for the agent must also be intelligible. As I explain below, there are primitive notions of action that do not have this requirement. But much of human agency is mediated by representational capacities, including (critically) rational capacities. Explanations of human action must therefore be *rationalizing explanations*.

Rationality, as I conceive it, is the capacity for conceptual thought, which is representation with propositional or conceptual content.⁴ *Propositional structure* is sentence-like,

⁴ The modifier “conceptual” is redundant, since I define thought as conceptual representational capacities, processes, and states. But I sometimes include the modifier anyway, to remind the reader of that usage, given how many different usages abound for the term “thought.” I say “propositional or conceptual content” because I think some mental states can apply conceptual attributives but lack propositional (subject-predicate) structure. Concepts

subject-predicate structure. *Concepts* are the constituents of that structure that give it content, and *inferences* are truth-preserving transformations between propositional or conceptual representations. Reasons, in the internal sense I mean, are conceptual representations that provide warrant, or function as a basis, for other conceptual representations, including evaluative and action-guiding conceptual states (conceptual conative states).⁵ The capacity for conceptual representation and inference mark an important border within representational psychology. Unlike lower level representational capacities like perception, thought includes a capacity for *pure attribution*, where properties, kinds, and relations can be represented as such without being instantiated in context-bound applications to particulars (Burge 2010a: 541–544). For example, in order to perceive sphericity, one must perceive some particular spherical object. But conceptual thought makes it possible to think about spheres or sphericity in general without representing any particular sphere.

Importantly, rationality (the mere capacity for conceptual thought) as such does not imply any metarepresentational or metapsychological abilities. That is, one can be able to form propositionally structured representations with conceptual content, but lack any concepts that refer to representational or psychological states as such. Conceptual thought is a constitutive precondition for metarepresentation, however, and metarepresentation is a constitutive

function constitutively to be elements in propositional representations. That is, concepts are what they are by virtue of being able to be constituents in predicates. Concepts have a kind of content that is apt to be applied in propositional structures. Therefore, I think that one cannot have conceptual representational capacities without having propositional representational capacities. Nonetheless, conceptual attributives can also be applied directly in representations that lack subject-predicate structure. For example, if I tell someone to think about elephants, they can bring the concept elephant to mind without forming any propositions about elephants. In addition, one can identify an object in one's environment as an elephant, applying the concept elephant, directly on the basis of visual perceptions alone. This application of the concept also seems not to involve any kind of subject-predicate structure.

⁵ Much ink has been spilled over the question of what reasons are. See, for example, Williams (1979), Korsgaard (1997), Richardson (1997), and Döring (2006). One basic distinction contrasts reasons as mental states with reasons as states of affairs. In this work I am principally concerned with reasons as mental states—this is what I mean by the “internal sense” of a reason, though I generally avoid the language of internal and external reasons because of the multifarious and inconsistent ways to which philosophers have employed those terms. In general, when I discuss reasons, context should make my meaning clear.

precondition for *critical reason*. Critical reason is the capacity to represent and evaluate reasons as such. Critical reasoners, like humans, can represent not just representational states as such, but can also represent the standards that constitute and apply to those representational states. Thus, critical reasoners can evaluate representational states (and actions) according to their constitutive norms. For example, thinking about a belief as a belief is a metarepresentational capacity, but thinking about truth and warrant is a critically rational capacity. Many philosophers use the term “rationality” to mean critical reason, which is understandable given that they typically mean *human* rationality. But I think it is a mistake not to distinguish clearly between conceptual thought and critical reason, since there are likely non-human animals (and also very young human children) with capacities for thought but not critical reason. And even within mature human cognition, much representation occurs at the level of thought but does not rise to the level of critical reasoning, just as much representation occurs at the level of perception but does not rise to the level of thought. The distinction is even more important when considering evaluative representations because critical reason is necessary for reflecting on values as such.

Understanding evaluative representations depends upon an understanding of *action-guiding representation systems*. I follow Burge (2022: 508) in calling action-guiding representations *conations* or *conative states*, which include both goal-setting representations and *commands*.⁶ Conative states, then, guide physical actions by setting objects in the environment as targets, selecting movement patterns, and executing and adjusting movement in real time (508–509). For example, catching a ball that has been thrown to me involves motor commands to raise my arm and close my fingers, and the particular movements will be guided in real time by

⁶ Burge (2022) focuses on *motor* commands, which are the most primitive bodily action-causing representational states, and *attention* commands, which govern primitive mental actions. I focus on the broader notion of a command that includes any representation functioning to cause and guide actions and activities, including higher level representations, such as intentions.

proprioceptive representations of where my arm shall go, along with perceptions of the ball's location and movement (506). Importantly, commands have veridicality conditions fulfilled by the execution of the action mapped out by the command.⁷ Thus, the motor command to raise my right arm is *made veridical* if it causes my arm to be raised in the way represented by the command. Similar points apply to more abstract (higher level) commands. Intentions are higher-level (conceptual) commands that function to organize and guide other conations. Intentions are constitutively rational states since they are conceptual.

Many evaluative representations are products of subrational processes and capacities.

The third claim expands the notion of evaluative representation beyond the scope of constitutively rational mental states. Any individual with rational capacities has capacities for belief, judgment, and inference, as those capacities are constitutive of rationality. But rational capacities are integrated with other more primitive representational capacities, including perception and emotion.⁸ *Perception* is a basic capacity for sensory representation, as it takes sensory information registration as input. But perception is marked by perceptual constancy, which is the fixity of representational content despite significant variation in proximal stimuli.⁹ For example, visual perceptions of a green sphere can refer to the same object (the sphere) and attribute to it the same properties (greenness, sphericity, approximate size, and so forth) even as the relevant optical sensory information changes (as when the lighting conditions and position of the sphere relative to the perceiver change). Perception is a relatively primitive capacity, prevalent throughout the animal kingdom, including among insects.¹⁰ There are also *intermodal*

⁷ For discussion and elaboration, see Burge (2010a: 337–341).

⁸ By “subrational” and “more primitive” I mean that capacities like perception are both phylogenetically and ontogenetically prior to the development of conceptual representational capacities.

⁹ See Burge (2010a: 408–413) and (2022: 64–105).

¹⁰ See Freas and Spetch (2023) for an overview of the representational capabilities of insects. See also Burge (2010a: 291–366, 514–518).

representations that receive inputs from different representational systems.¹¹ I hypothesize that some of these are perception-like representational capacities that facilitate interaction between rational and subrational capacities. Thus, I think there are evaluative representations with conceptual content—or at least, content that is sensitive to conceptualizations—but perceptual presentational structure. I develop this point more in the next chapter, but for now just note that I think these perceptually formatted intermodal representations, when they have the right sort of evaluative content, are good candidates for what Brewer (2009) calls *desires*, which I discuss below.

Highlighting the role of subrational capacities, like perception, matters because the strongest version of an evaluative outlook account will avoid over-intellectualizing agency, even for rational and critically rational agents. Many actions are mediated by representation even if they are not rational, and it is unlikely that all, or even most, rational explanations of action primarily involve explicit evaluative judgments or evaluative inferences. Such a requirement would not only make commonplace cases of akrasia difficult to explain, it would exclude the possibility of providing rational explanations of the actions of young children and non-human animals. And young children and many non-human animals seem, pretheoretically at least, to act for reasons even when they do not make explicit judgments about value or about what they have reason to do.

At least in humans, evaluative representations function to be elements in an overall evaluative understanding, or conception of the good. The fourth claim clarifies how evaluative representations issuing from subrational processes can nonetheless be elements in successful rational explanations. Essentially, it claims that rational capacities function to exercise control

¹¹ See Burge (2010b) for some development of this point. See also Carey (2009).

over subrational capacities. In critically rational agents like adult humans, a general evaluative understanding functions to guide all evaluative representations by being the framework in which various evaluative representational capacities are integrated. The relation between evaluative understanding and the various subrational evaluative representations should be understood as a two-way functional relationship. In general, subrational representational capacities like perception function to inform and shape conceptual representations, but conceptual frameworks also function to influence (many) subrational processes. These functions are rational functions. There are natural norms that govern when conceptual representations like beliefs should and should not be revised in light of lower-level representations, like perceptions. These norms are constitutive aspects of the conceptual representational capacities themselves, and thus they are constitutive of rationality. It is not constitutive of perception that it function to shape understanding, but it is constitutive of understanding that it functionally incorporates and influences perceptual capacities. Thus, understanding guides all representational capacities in the sense that all representational capacities have a role to play in the larger centralized representational system. And there are fundamentally rational norms that govern how various representational capacities fit into the overall system.¹²

Evaluative representations by nature have motivational efficacy. This fifth and final claim insists that evaluative representations by nature have a motivational function. Indeed, what it is to act rationally is to be moved by evaluative representations informed by an overall evaluative understanding. Thus, the evaluative outlook model explicitly rejects Humean and non-cognitive views according to which motivating mental states are not representations.

¹² In individuals without rational capacities, there are other, more primitive, ways that representations are coordinated within the individual's overall psychology. In my view, understanding is the capacity that coordinates representational capacities in rational individuals. I discuss understanding at length in Chapter 4.

To get a grip on the motivational efficacy of evaluative representations, keep in mind that evaluative representations as I conceive them are a type of conative state. Conceptual conative states include *intentions* and *evaluations*. Intentions are the conceptual equivalent of lower-level action commands. Evaluations function to set goals for intentions. Evaluation-like states occur at the level of perception as well, and include attributions of perceptual categories like food and danger.¹³ These attributions are evaluation-like in that they function constitutively to set goals for action. For example, the attributive food functions to serve as the goal of the perceptual-level action command to eat (and other related action commands). At the level of thought, evaluations become more complex and are sensitive to the individual's overall evaluative understanding, just as non-evaluative conceptualizations are sensitive to one's understanding in the relevant domain. At the level of thought, evaluative representations have the same functional structure: conceptual evaluative attributives function to be goals in conceptual action commands (intentions).

These five theses are abstract and allow for a variety of evaluative outlook accounts. But before I look at one such account in detail, I want to pause and reflect on the nature of value itself. Much of this next section merely defines the terms I use to discuss value, though in doing so I lay out some important distinctions. Readers more interested in the practical psychology aspect of evaluative outlooks may wish to skip or merely skim this material.

III. On the Nature of Value

I use the term *value* as the most general term for practically significant or teleologically relevant particulars, properties, kinds, relations, and states of affairs. If something is practically significant, then actions and attitudes ought (absent defeaters) to be responsive to it; it constitutes

¹³ Burge (2010a: 323–325) calls these perceptual categories “functional attributives.” I discuss perceptual level conative states in more detail in Chapter 2.

a practical standard. In my terminology, if something is practically significant, it *matters*. The notion of practical significance marks a general kind of value, though usually these values are *not* explained at the most fundamental level by their practical significance. For reasons that will become clear, I reject response dependent accounts of value according to which all values just are (or are dependent upon) standards governing actions and attitudes. It is true that some values are by their nature practically significant, but in general a valuable thing is not valuable in virtue of its practical significance; rather, it is practically significant in virtue of being valuable. Here, as elsewhere, the order of explanation is important. Moreover, some values are not practically significant.

Some values are *goods*, a term I reserve for relative values. Values can be relative in two general ways. First, some values are relative because they are teleologically relevant. If something is teleologically relevant, it partly constitutes the *success* of some function, the *excellence* of some kind, or the *flourishing* of some individual. Success is goodness relative to function. For example, “a good heart pumps blood,” and “it is good for the tree to get enough water.” Excellence refers to a particularly high level of functional success, often determined by an ideal toward which a functional system is directed by virtue of the kind of system that it is. (Virtue and perfection are kinds of excellence.) For example, a tree might have an excellent root structure, while another tree has a comparatively poor one, even though both root structures meet the minimum standard for functional success (they keep the trees alive). Flourishing is a kind of functional success specific to living individuals. There are biological kinds of flourishing as well as psychological kinds. Flourishing has different dimensions depending on the capacities and nature of the individual in question. Human flourishing includes things like biological and psychological health, moral virtue, and happiness, for example.

The second general kind of relative value is the good of conscious individuals. These values are relative to the *well-being* of individuals. Well-being, as I use the term, is not a constitutively teleological notion. Something like pleasure, for example, might be good for an individual regardless of whether it fulfills any function or purpose. It is likely, of course, that goods of well-being are always teleologically relevant because all individuals with well-being are agents, who thereby ought to pursue their own well-being, at least within reasonable limits.¹⁴ As a matter of fact, then, all well-being is practically significant from at least one perspective (see below). Moreover, the well-being of an individual contributes to that individual's flourishing—the flourishing of conscious individuals is partly constituted by their well-being. But whether something is an aspect of the well-being of an individual does not fundamentally depend on functional success, excellence, or flourishing.

Success by definition is always functionally relevant, but it is practically significant only hypothetically. That is, the success of a function only ought to constrain an individual's actions if either the individual correctly aims at the success of that function, or the function itself constitutes the individual's practical success. For example, in general survival partly constitutes the practical success of every living individual, so an individual's biological function of nutrition is a practically significant constraint on that individual's actions. On the other hand, it is good for a tree to get enough water, but that only matters if the biological functioning of the tree matters independently, and the biological functioning of the tree does not matter merely in virtue of its being functional. Similarly, it is good for a car to receive regular oil changes, but that only

¹⁴ Some might apply the term well-being to classes of entities other than agents, such as living things. But I restrict the term to the good of conscious individuals, who as far as we know all happen to be agents, at least in a primitive sense. This primitive sense of agency, which is not the most primitive sense, just indicates behavior guided by representational capacities. Of course, if it turns out that there are organisms who have consciousness but lack representational capacities, then it is false that all individuals with well-being are agents in this sense.

matters if the well-functioning of the car matters, which it may not. This contingency is why not all values are practically significant. Some goods don't matter.

Worth is non-relative value. Worth matters from any perspective capable of recognizing or responding to it. It is part of the nature of worth that it necessarily matters, but the nature of worth does not wholly consist in its practical significance. We may be able to pick out entities of worth by pointing to their necessary practical significance and the appropriateness of certain attitudes and actions toward them. But an entity's worth does not consist in this appropriateness; the nature of worth is not wholly constituted by or reducible to the fact that certain responses to it are correct.

Worth is not goodness. In itself, worth is not teleological. For example, human persons have worth, but the worth of human persons is not constituted by the fulfillment of any function or purpose. Of course, there is a minimum threshold of functional success necessary for an individual in fact to be a person. But once these minimum conditions are met, the worth an individual has in virtue of being a person does not depend on any further functional success. Nor does the worth of persons consist in their excellence or well-being (or the excellence or well-being of anything else). It is true that because human persons have worth, their excellence and well-being matter; however, the order of explanation is clearly from the worth of persons to their good. That is, the good of human persons matters because human persons matter. I think this claim is warranted by reflection on our moral practices, on attitudes like admiration, respect, and love, and on the role of happiness in the good life. I develop these points in greater detail in later chapters.

For an individual with a psychology, and thus with a perspective, the individual's own excellence and well-being typically constitute practical constraints for that individual. For

example, it is good for the rabbit to avoid being eaten by the hawk. There is at least one perspective from which the rabbit's continued existence and well-being are practically significant—namely, the rabbit's own perspective. The rabbit *should* treat its own continued existence and well-being as practically significant, where “should” indicates a natural norm (see below). This point does not imply that rabbits represent their own well-being or survival explicitly. They may or may not (I suspect they do not), and whether they do is an empirical question about the psychological capacities of rabbits. A practical constraint is a norm governing the range of actions an agent takes or goals an agent pursues. Such a norm might depend on the agent explicitly representing something as valuable, but it need not.

A *norm* is a standard that specifies the minimum fulfillment of a function, a threshold of excellence, or an aspect of well-being. A *natural norm* is a norm that applies independently of its being appreciated, accepted, set, or represented in any individual's psychology or within any group's collective practices or attitudes (Burge 2010a: 311–312). That a heart pump blood minimally well, that a visual representation be veridical, that a belief be true, and that a human person avoid wrongdoing are all natural norms. They apply regardless of whether they are accepted, aimed at, or explicitly part of any social practices. That is, there are entities that can be assessed according to such standards merely in virtue of being the kinds of things that they are, and having the natures that they have.¹⁵

¹⁵ Moral norms are something of a special case since they depend upon individuals' being capable of representing reasons as such. That is, moral norms apply to all and only critically rational agents, which are reasons-responsive agents that can represent and evaluate the sufficiency of their own reasons. Moral norms apply to individuals and groups capable of recognizing them. Nonetheless, a moral norm's applicability does not depend on its actually being recognized by the individuals or groups to whom it applies. That an agent or group fails to recognize a moral norm does not entail that the norm does not apply to that agent or to individuals in that group.

Worth implies a categorical practical norm.¹⁶ The worth of human persons should guide the actions and goals of all agents capable of recognizing their value. This is a natural norm, and it helps explain why moral norms are natural norms.¹⁷ The only condition on the practical significance of worth for any individual agent is that the agent be capable of recognizing or responding to it. If there were no agents capable of recognizing worth, worth would still exist and have the nature that it has.¹⁸ The worth of persons is also *independent*, since it is grounded in the nature of human persons themselves. In addition, the good of human persons is a *final good*, since its practical significance is not grounded in some other more fundamental good that it produces, causes, or constitutes.¹⁹

The mark of a final good is that it is appropriate to value it for its own sake, or as an end in itself. A non-final good is appropriately valued for the sake of something else, but not for its own sake. Among non-final goods, some are *merely instrumental*, meaning that they are practically significant only in virtue of further goods that they causally produce, and some are *merely constitutive*, meaning that they are practically significant only in virtue of further goods that they help constitute, usually by being necessary components or instantiations of those further

¹⁶ My use of the terms “categorical” and “hypothetical” derives, of course, from Kant (2002). As I use the terms, they apply to practical norms. A norm is just a minimally adequate standard of functional fulfillment (Burge 2010a: 311). A practical norm, then, is a minimum standard of success for an agential function. Such a norm is categorical if it applies to an individual’s actions by virtue of the individual’s naturally necessary aims—the aims the individual has (or should have) by virtue of the individual’s nature as the kind of entity that it is (its membership in a particular category). A practical norm is hypothetical if it applies to an individual’s actions merely by virtue of the individual’s contingent aims—aims which could be otherwise than they are.

¹⁷ See Burge (2010a: 314, fn. 36). I recognize that the claim that moral norms are natural in the sense I describe is controversial. It is certainly true that many specific social norms that we naturally think of as moral are not natural in this sense. Nonetheless, I think that at a certain level of abstraction, basic moral norms deriving from the worth of individual persons are natural; they apply to all agents capable of recognizing them, regardless of whether particular agents do in fact recognize them, accept them, or aim at them.

¹⁸ In fact, this likely was the case for long periods of evolutionary history. I think that all conscious individual subjects have worth, and there were almost certainly conscious individual subjects before agents (human beings) evolved who were capable of recognizing and responding appropriately to the worth of those subjects.

¹⁹ This does not mean that it is always correct to pursue or promote every person’s good; that would be impossible. Rather, the point is that each person’s good matters by virtue of their worth as a human person.

goods. Of course, some final goods should be valued for their own sakes *and* for the sake of something else, either because they also have instrumental value or because they constitute some further good.

Worth creates practical constraints in at least two ways: positively and negatively. Worth creates negative constraints in the sense that it limits which actions and goals are reasonable or permissible for agents capable of responding to it. Thus, worth governs rightness and wrongness. Worth also positively guides agents capable of appreciating or recognizing it towards various activities and goals that matter. Thus, I call the activities of such an agent who is living well *choiceworthy*. (In Chapter 5, I consider the choiceworthiness of activities, including the question of what makes human life itself choiceworthy.)

The relation between goodness and worth presents an important philosophical puzzle, which I hope to shed light on in the course of this dissertation. Here I merely make some preliminary remarks. First, the three most obvious candidates for kinds of entity that have worth by nature are all teleologically structured beings with their own good as individuals of their kind: human persons, sentient beings, and living things. I take it to be obvious and uncontroversial that human persons have worth. The value of human persons is manifest in the fact that it matters what happens to them; it matters whether or not they are virtuous; their well-being matters; in short, human persons matter. Similar points apply to sentient beings in general, and I speculate that consciousness as a capacity grounds the worth of individuals who have it. But this view is not uncontroversial.²⁰ Even more controversial is the view that living things as such have worth

²⁰ See Korsgaard (2018) for discussion of the Kantian view that non-human animals do not matter for their own sakes. Korsgaard's own view attempts to bring a direct duty to value non-human animals for their own sakes under the constructivist umbrella, grounding the value of non-human animals in a purely formal requirement that non-human animals be valued if we are to value anything at all (131–145). I do not find such constructivist accounts compelling, and I offer some general critiques of Korsgaard's constructivism in Chapter 5.

(I suspect that they do not, but I shall not take up the issue here). For my purposes, I begin with human persons as paradigms of worth. Moreover, since some goods are available only to persons (because of their capacity to think critically about their own psychological states and well-being), focusing on persons can help flesh out a complete taxonomy of goods. For now, I simply note that I do not think human persons are the only things that have worth, and I set aside the question of what grounds the worth of entities who matter in themselves.

Now, having outlined these distinctions within the realm of value, I return to evaluative outlooks. Specifically, I look at how Brewer (2009) motivates his particular evaluative outlook account, and I add two additional claims to the five key commitments of the evaluative outlook model discussed above. These additional claims aim to capture Brewer's insights into the importance of phenomenology in mature evaluative understanding.

IV. Dialectical Activity, Desire, and Phenomenal Consciousness

Talbot Brewer's (2009) account of practical reason begins with *dialectical activity*. According to Brewer, dialectical activity is temporally extended activity carried forward by an ongoing understanding of the activity's value (37). This value includes the goods or excellences that are internal to, and partly constitutive of, the activity. A dialectical activity has a complexity, a depth, that allows engagement with it to be carried forward over time by increased understanding of the activity's value. That understanding is not a static achievement, but a running process developed by one's engagement in the activity itself. Thus, the way an agent engages in a dialectical activity can change radically over time as her appreciation for and understanding of what she is doing sharpens and deepens. This reiterative process of engagement

with an activity, refinement of understanding of the activity's value and point, and then reengagement with the activity, can be extended indefinitely—hence the label, *dialectical*.²¹

Dialectical activities are central to human life.²² Taking them seriously places important constraints on any plausible account of practical reason. Brewer's suggestion for incorporating dialectical activity into an account of practical reason centers on a particular notion of desire, which Brewer labels the *evaluative outlook* account of desire (2009: 24–32). On this model, desires are “apprehensions of value” that function in our practical thinking much like perceptions function in theoretical thought. Just as perceptions are “conceptually structured takes on” our physical environment, desires are conceptually structured takes on our evaluative circumstances (25).²³ Brewer argues for this account of desire largely on the grounds that it is the best explanation for the role that desires play in rational explanations of action. In order to make sense of any human action *as rationally evaluable action* (and not merely as behavior), Brewer claims that we must have some sense of what *good* or *point* the agent saw in the action (27–29). Desires, he argues, can play this role if they are apprehensions of value: perception-like evaluative representations.

Engagement in dialectical activities, at least in successful cases, deepens the agent's understanding of the values that make those activities choiceworthy. These values are constitutive elements of the activity. By understanding these values, the agent comes to

²¹ See Brewer (2009: 37–67) for an elaboration of the notion of dialectical activity.

²² There is a question of how central a place to afford dialectical activity in an account of practical reason. Cordner (2014), for example, argues that Brewer overstates the usefulness of the notion for capturing core elements of rational life. Nonetheless, even if dialectical activity does not present a complete picture of inherently valuable activity, I think it does present an important piece of the puzzle. It is enough for my arguments that dialectical activity is centrally important to human life; it need not be the only important kind of activity. That said, I think that Brewer's account has room to accommodate Cordner's critiques.

²³ Perceptions, as I understand them, are not conceptually structured. Brewer's notion of perception is more capacious than mine. The claim that desires are conceptually structured needs qualification, too, but for now we can accept this characterization as Brewer presents it. I discuss these issues in Chapter 2.

understand better the final value of the activity itself. For example, an artistic endeavor like musical composition has a value that cannot be understood without grasping the particulars of that activity, including what counts as excellent musical composition. To understand the value of musical composition is to see the various actions and activities that make up musical composition as worthwhile, and pursuing this activity requires developing an ongoing understanding of that value.²⁴

If Brewer is correct that explaining this kind of activity is indeed central to an account of practical reason, then we must explain (i) what values are internal to these sorts of activities, and (ii) what it means to deepen one's understanding of those values. In considering these questions, two points stand out right away. First, understanding the value of an activity partly depends upon understanding the activity's proper place within a good human life (Brewer 2009: 40–41).²⁵ So human well-being or flourishing is connected to a proper understanding of dialectical activity. Second, understanding the value of an activity must also involve a grasp of how to pursue the activity excellently. Understanding the value of an activity, then, must involve some grasp of its constitutive ideals and aims (43–45). These two points give some sense of the content of the

²⁴ We must take care not to confuse two nearby claims here. Brewer's account claims that engaging in dialectical activity requires and is carried forward by an evolving understanding of the activity's purpose and value. This is *not* the same as saying that the agent *aims at* (either explicitly or implicitly) increased understanding of the point of her activities. If I am engaged in the activity of climbing a mountain, for example, I am not aiming at improving my understanding of the value of mountain climbing. I am probably developing that understanding as I go about the activity, but my purpose is to climb the mountain. My attention will be on details of the task at hand—ascending that rock face, traversing this ice flow, and so on—and perhaps on features of the experience—the beauty of some vista, or the satisfaction of conquering a challenging obstacle. My attention is decidedly not on my own understanding.

²⁵ Possibly, the agent must come to understand the nature and value of an independently valuable entity around which the activity centers. For example, activities of love revolve around the beloved, a being whose value is entirely independent of the activities of love, and whose value grounds the value of those activities. Love and friendship are central examples of dialectical activity that Brewer develops in detail (see especially 2009: 236–285). I am not certain whether all dialectical activities must somehow center on such an independently valuable entity; I suspect not. It does seem likely, however, that all dialectical activities connect in some way with value that has a source independent from the perspective of any individual agent. Nagel (1986) and Wolf (2012) make similar points about the perspective-independence of certain kinds of value. I discuss Wolf's view in detail in Chapter 5.

relevant understanding. But the presentation of that content is also important. For Brewer, understanding the value of dialectical activities is closely connected to phenomenal conscious experience. He does not develop this point in great detail, but I think it is essential to his view. In what follows I build on Brewer's insights about the role of phenomenal consciousness in evaluative outlooks, and I present two additional theses that incorporate phenomenal consciousness into the general evaluative outlook model.

Brewer's evaluative outlook view is motivated in large part by dissatisfaction with a view that he calls "propositionalism." Propositionalism, as Brewer presents it, tries to explain action without invoking any irreducibly evaluative representational content (2009: 15). It is this propositionalist ambition that leads to the common belief-desire model of rational action, according to which all rational actions can be explained as the pairing of a desire that some state of affairs be the case and a belief that the action in question will make that state of affairs the case.²⁶ But Brewer thinks that reflection on dialectical activities and their central role in human life gives us reason to reject propositionalism.

Now, I am less concerned with the argument against propositionalism than Brewer is, but his reasoning here highlights the importance of phenomenal experience in evaluative understanding. Appreciating an activity's place within a good human life and grasping how to pursue an activity well are cognitive achievements manifested in the avenues for action that an agent sees spread out before her. One might explain this in terms of reasons: what the agent sees herself as *having reason to do*. But it might also be less intellectualized than that: what the agent is *drawn to*. This second formulation points to the motivational efficacy of the kind of evaluative understanding that Brewer focuses on. It is not a detached, abstract understanding. It cannot be

²⁶ See Davidson (1980) for the classic exposition of this view.

adequately explained by listing propositions that the subject affirms. There is an important phenomenal dimension to the way that evaluative understanding moves a person to act, and this qualitative phenomenology of finding something appealing explains the difference between judging that one has reason to do x and genuinely desiring to do x . Brewer describes this difference as the difference between desiring engagement in the activity and merely desiring *that it be the case that* I engage in the activity (45–46). These genuine desires are products of our evaluative understanding in that they are “structured by the [evaluative] concepts we have mastered” (31). And mature evaluative understanding must incorporate the qualitative phenomenal richness of our experience, since this phenomenal quality partly constitutes our desires as the kind of mental state that they are, and our evaluative understanding functions to shape those very desires. Moreover, our evaluative understanding of an activity must be capable of being shaped by our qualitative experiences of engaging with the activity. Thus, in order to grasp the value of an activity, one must understand the goods internal to and partly constitutive of the activity, and this understanding in its ideal form issues in an essentially qualitative phenomenal appreciation of the activity’s value.

Phenomenology is a feature of *consciousness*. Consciousness is the subjective presentation of mental content by and to the individual.²⁷ Upon reflection, consciousness as such does not seem to be constituted by any kind of about-ness or representational content. That a state is subjectively presented in the individual’s mind is sufficient for that state to be conscious, even if the state does not function to represent anything at all. A basic example is pain. A pain state is a conscious state if the individual feels the pain, regardless of whether the state has any correctness conditions. Of course, many conscious states are in fact representational. For

²⁷ See Burge (2007a, 2007b) for discussion of this way of explaining consciousness. See also Chapter 3 of this dissertation.

example, conscious perceptions and conscious beliefs are both conscious and representational. Once again—and this point is methodologically fundamental—constitutive conditions concern what makes a thing the kind of thing that it is, or what explains a thing’s having the nature that it has. The nature of consciousness is subjective presentation, and we know this because we can pick out conscious states, like pain states, by their subjective presentation alone. The conscious nature of a state is not called into question if the state has no representational content. But a state’s purported conscious nature is undermined if it has no subjective presentation.²⁸ It may be that pain states sometimes function to represent bodily damage, for example, and they certainly can function to *indicate* bodily damage; but their nature as conscious states does not hang on whether or not they have such functions. These considerations lead to the conclusion that consciousness as such is not constitutively representational, nor is it constitutively functional.²⁹

Sentience is a basic kind of consciousness. Sentience, as I use the term, is a capacity for *phenomenal* conscious states. Phenomenology is *qualitative* content, which has essential features that cannot be captured in quantitative, mathematical, or purely formal terms. For example, conscious color perceptions (in humans and similar animals, at least) have qualitative modes of presentation. What it is like to see red has no fully quantitative or formal explanation.³⁰ It is difficult to explain the nature of qualitative content aside from the contrast with quantitative

²⁸ Importantly, a phenomenal state, like a pain state, need not be subjectively presented in current consciousness in order to be a phenomenal state. If there are unconscious pains, for example, the state’s nature as a pain state will depend upon what it is like to be in the state when it is conscious. Thus, there may be unconscious states that depend for their nature as the kind of state that they are on a capacity for phenomenal consciousness. See Burge (2007a, 2007b) for discussion of this point.

²⁹ See Burge (2007a and 2007b) for additional development of and argument for this point, along the same lines that I present here. See Block (2002) for the differentiation between phenomenal consciousness and representational (or “access”) consciousness, on which Burge builds.

³⁰ The “what it is like” characterization of conscious states comes, of course, from Nagel (1974). The particular characterization of this idea in terms of qualitative versus quantitative content is my own. Of course, perceptual experience of color does have important quantitative elements, since the physical phenomena—light spectra—that such perceptions function to represent have mathematical properties. Thus, the color spectrum can be explained in mathematical terms, but the qualitative experience itself of what it is like to see red cannot be so explained. See Chapter 3 of this dissertation for further development of this idea.

content. Another way to get at the idea is the contrast between form and substance. A purely quantitative property is entirely formal, or structural. A qualitative property is substantive, and cannot be fully explained by its structural features, if it has any.³¹

Some phenomenal states are *felt*. Felt states are bodily sensations of the individual. These states need not be representational, but they may be, and they are often accompanied by proprioceptive locations, as when one feels a tickle on the bottom of one's foot. Some felt states are *affective*, which means that they occur globally, either throughout the individual's whole body, or without being indexed to any specific bodily location. Among affective states, some have representational content, and these are *emotions*.³²

A conscious representation has a different kind of content than an unconscious representation, even if they both have the same correctness conditions. For example, a conscious perception of a red body is a different kind of psychological state than an unconscious perception of the same red body. Some representational states have a conscious *mode of presentation*. A mode of presentation is a way of presenting mental content, and different modes of presentation indicate different psychological kinds. So a conscious mode of presentation is a different psychological kind than an unconscious one, even when the two representations in question have the same veridicality conditions (Burge 2010a: 76). Indeed, there is some empirical evidence that conscious and unconscious representational states are functionally different.³³ Along similar

³¹ These epistemic points about what kinds of explanations are available are not definitional. For instance, when I say that qualitative properties cannot be explained mathematically, I am not defining what qualitative properties are in terms of the unavailability of mathematical explanations. Rather, I aim to evoke understanding through example and comparison. I do think that unavailability of a mathematical explanation is a necessary condition for qualitative content, but it is nonetheless not a definition since it does not positively explain the nature of qualitative content.

³² Burge (2022: 507–508, 546 fn. 771) briefly develops the notion of emotion as representational affect. I consider emotion, including Burge's conception, in detail in Chapter 3. Although I think all representational affective states are emotions, I am not certain that all emotions are representational affective states.

³³ See Birch (2022) for empirical arguments for the *facilitation hypothesis*, which claims that conscious states and unconscious states facilitate (or cause) different kinds of cognitive and behavioral tasks. One limitation of

lines, I speculate that a *phenomenally* conscious state is of a different psychological kind than a non-phenomenal conscious state, and that there may be functional differences between phenomenal states and non-phenomenal ones.

So when we interpret Brewer's claim that a desire is an essentially phenomenal appreciation of value, we should interpret that as a claim about a kind of motivating representational state. But it need not be the only kind of motivating representational state. Phenomenology is not constitutively necessary for motivation, though it does influence motivation. Phenomenal motivational states are functionally different from non-phenomenal ones. For example, at the level of perception, there are likely motivating representations with infinitival structure, such as to eat. These representations need not have any phenomenal content, but are nonetheless motivating conative states. Similarly, at the level of thought, there are intentions that have no constitutive phenomenal quality—perhaps there is nothing that it is like to intend to wash the dishes, and yet this intention can nonetheless move the agent to act.³⁴ Brewer's point, I think, is that there is an important class of motivating representational states that do have essentially phenomenal evaluative content. And these states, which he calls desires, are important for understanding human practical reason.

At this point, I think we can add two additional theses to our list of commitments of the evaluative outlook view. In addition to the first five theses listed above, Brewer's evaluative outlook view also insists upon the following:

these arguments is that the empirical evidence presented for the facilitation hypothesis does not show that it is consciousness as such that explains the difference between the causal powers of conscious versus unconscious states. It could be that some other psychological capacity, such as centralized cognition, explains the difference and also happens to covary with consciousness. For Birch's purposes, this limitation is no issue, since Birch does not aim to explain the nature of consciousness as such, but only to establish what kind of evidence would secure strong inferences that consciousness exists in various non-human animals. In order for some ability to be evidence of consciousness, it need not be that consciousness is constitutively necessary for that ability.

³⁴ I discuss motivation, especially rational motivation, in more detail in Chapter 4.

- 6) Some evaluative representations are partly constituted by their phenomenal quality.
- 7) The phenomenal quality of an evaluative representation (if it has phenomenal quality) partly explains its motivational efficacy.

Now, Brewer has in view a sophisticated conception of agency. He writes, “To bring an event under the concept of action is to view it... as something produced by a *conscious* being out of that being’s appreciation of reasons for producing it” (33, my emphasis). Brewer articulates this conception of agency to rebut Humean and non-cognitivist views of motivation, which claim that representational states by themselves cannot motivate. He is clearly focused on *human* agency. Notably he insists that human agency consists not just in representational capacities, or reasoning capacities, but also *consciousness*. His point, I take it, is that we cannot explain human actions without reference to how human beings *consciously* apprehend reasons to behave in various ways. I think, however, that considering more primitive agents can be helpful here, not just to deflate Humeanism about motivation, but also to emphasize the role of qualitative conscious states in human thought and action.

Agency as such is very primitive. In its simplest form, *animacy*, agency is simply behavior. Behavior is functional movement by the whole entity. Thus, many simple organisms exhibit primitive agency of this sort. Animacy requires only functional organization and whole-entity movements coordinated or caused by functional capacities of the individual.³⁵ Indeed, many human-made artifacts are primitive agents in this basic sense. This notion casts a wide net. When agency is centrally coordinated and guided by representations, we have psychologically-mediated agency. (This is not to be confused with agency *in* psychology, as

³⁵ See Burge (2010a: 326–334) for discussion of the notion of primitive agency.

when one exercises control over one's thoughts.) All organisms with perceptual capacities exhibit this type of agency. Once there are representations that coordinate whole-entity movement and that have propositional structure and conceptual content, we have rational agency. Reasons are just representations that function to guide thought or action and that have propositional structure or conceptual content. Finally, critically rational agency includes reflective capacities for representing and evaluating the reasons that the agent acts on.

The notion of agency Brewer articulates in the quotation above might be restricted to critically rational agents or merely rational ones, depending on what "appreciation" means. If it means representations of reasons *as reasons*, then he has in view only critically rational agents. But if "appreciation of reasons" just means responsiveness to reasons, then less sophisticated thinkers who have capacities for propositionally structured thought but who lack critical reason will be included. With his focus on human agency, Brewer does not pursue this distinction; after all, typical human beings either have or are developing capacities for critical reason. Nonetheless, the distinction is important more broadly, since it bears on how far our rationalizing explanations of actions can be applied to non-human agents, as well as to human agents who lack some critically rational capacities. Moreover, many actions of mature human beings occur at a lower level, not directly mediated at all by reasons, much less critical reasoning—various habitual behaviors, for example.

Keeping an eye on primitive notions of agency suggests something else as well. Brewer's requirement that rational agents be conscious cannot be derived merely from reflection on the concept of action. Rather, I think this requirement comes from reflection on *paradigmatic rationalizing explanations* of action. Brewer thinks that such explanations make essential reference to phenomenally vivid apprehensions of value, and so he concludes that (rational)

agency must necessarily involve a capacity for phenomenally conscious states. But some care is required here. I think it is probably true that all thinking creatures (that is, all creatures capable of conceptual and propositional representations) are conscious. It may even be *necessary* that all thinking creatures are conscious. But it is not clear whether this necessity is due to the nature of rational thought. That is, it is not clear that consciousness is a constitutive element of reason as such. If it is, that would be a major discovery. In any case, the role of phenomenal consciousness in rational thought and action is a worthy target for inquiry. If Brewer is right that the conscious phenomenal quality of representations plays a key role in rational motivation, that, too, would be a significant contribution.

These are not matters that can be settled from the armchair, but I hope some progress can be made by reflecting on representation, action, motivation, reason, and value, among other things. Thus, I hope the reader will appreciate the exploratory spirit in which I develop and argue for the evaluative outlook view.

The first point to make is about the place of dialectical activity in a proper conception of rationality. My initial gloss on the notion of rationality presented it as a capacity for representations that have a certain kind of structure: *propositional* structure.³⁶ Recall that the key capacity that marks propositional representation is inference: truth-preserving transformations between propositional (or conceptual) representations, where truth is just the kind of veridicality (correctness) proper to propositional representations. Emphasizing dialectical activity, however, puts pressure on this largely formal conception of rationality, since making sense of dialectical

³⁶ The jargon is somewhat regrettable, but it follows typical philosophical usage, and Burge (2010a, 2010c, 2022) in particular. I might have said *sentence-like* structure instead, since the essence of propositional form is subject-predicate structure.

activities seems to require giving a central place to a certain kind of rational mental *content*.³⁷ If Brewer is on the right track, then the structure of dialectical activity mirrors the structure of human practical reason. But we cannot make sense of that structure without positing evaluative representations with distinctively phenomenal character. On Brewer's view, we are drawn to dialectical activities by a phenomenally vivid, yet conceptually inchoate or indistinct, sense of the value of those activities or their objects. Nonetheless, these phenomenally vivid evaluative representations are subject to influence by the understanding, since engagement in the activity develops our conceptual grasp of the activity's value, which in turn shapes the phenomenally vivid evaluative representations that draw us into further stretches of engagement.

These phenomenally vivid representations, which Brewer calls "desires," are one of the primary targets of my analysis. I hope to show that they are constitutive elements of rationality, even if they are not themselves constitutively rational states. Rationality, as I understand it, depends upon other mental capacities, including phenomenal consciousness, that are not themselves constitutively rational capacities, but that become incorporated into, and functionally governed by, rational processes.

Another point is to emphasize how little we understand consciousness. Some of the relevant questions are empirical. I do not think we know, for example, whether consciousness is more primitive than representation. Settling empirical questions about when consciousness actually developed as a capacity in evolutionary history will shed some light on the nature of consciousness. But our understanding of the nature of consciousness is limited enough that the empirical questions are very hard to investigate. It is hard to see clearly how they would be

³⁷ To proliferate even more jargon, we need a semantic account of rational mental states, not just a syntactic one.

answered in our current state of understanding.³⁸ Reflection on what little we do understand about consciousness, however, yields a few speculative hypotheses.

First, as I indicated above, the nature of consciousness is not constitutively functional.³⁹ That is, it is not by virtue of being a functional state or fulfilling a function that a conscious state is conscious. A conscious state is conscious by virtue of being presented in subjective awareness. To reiterate my previous example, a pain state is a pain state by virtue of its phenomenal quality, regardless of the function of pain. This point would in principle hold even if it turns out that all conscious states have functional roles. Second, phenomenal consciousness is more fundamental than non-phenomenal consciousness.⁴⁰ That is, non-phenomenal conscious states depend for their conscious presentation on a general capacity for phenomenal consciousness, even if they are not themselves constituted by any phenomenal conscious qualities. For example, the conscious thought 5 is a number is not constitutively phenomenal. Such a thought is the thought that it is by virtue of having the propositional content that it has, and it would be the same thought even if what it is like to have that thought varies wildly. You and I can both have the conscious thought 5 is a number and therefore have the same thought, even if the qualitative presentation of that propositional content is very different in your mind and in mine (or even if one of us lacked phenomenal consciousness altogether). By contrast, if you and I both consciously see the color red, our mental states must both, in order to be instances of consciously seeing red, have a certain phenomenal quality. As a kind of mental state, consciously seeing red is constitutively phenomenal. Thus, some conscious states are constitutively phenomenal and others are not.

Nonetheless, I think non-phenomenal conscious states are parasitic on a capacity for phenomenal

³⁸ Nagel (1974) makes a similar point. Notably, it is as relevant now as it was then.

³⁹ See Block (2002) and Burge (2007a) and (2007b) for further articulation of and argument for this point.

⁴⁰ See Burge (2007a) and (2007b) for some tentative arguments to this effect. I share Burge's tentativeness about this point, but it seems to me correct nonetheless.

consciousness. Whenever mental content is presented in subjective awareness, it takes on some kind of quality—there is something that it is like to be in that particular (token) state—even if that qualitative element is not constitutive of the state as the kind of mental state that it is.

Both of these speculations are supported by reflection on our scientific and common sense understanding of consciousness. I think they are warranted a priori. But they are both defeasible. It is possible, at least in principle, that consciousness will be reduced in a future scientific theory to some particular complex functional kind, for example. In the absence of such a theory, however, I think both of these points about the nature of consciousness are warranted, and they ought to serve as a basis for our theorizing about consciousness, both philosophical and scientific. So whatever role consciousness and phenomenology play in a correct theory of practical reason, other things being equal, such a theory ought to account for, or at least respect, the non-functional nature of consciousness itself, and the fundamentality of phenomenal consciousness. An evaluative outlook model can, I hope, satisfy this criterion.

V. The Way Ahead

The chapters that follow aim to lay some of the psychological groundwork for a viable and persuasive evaluative outlook account of practical reason. The territory is vast. I cannot hope to cover it all. But if I can sketch the borders and some important landmarks, I will count myself successful.

In the second chapter, I begin with the notion of representation. I sketch an account of perception, based on Burge's (2010a, 2010b, 2022) philosophical analysis of the relevant scientific literature, and contrast perception with more sophisticated representational capacities. Then I consider how representations, especially perceptions, could have evaluative content. Of

course, some kinds of evaluative content are more relevant to theorizing about practical reason than others. I conclude that the kinds of evaluative representation of most interest to the evaluative outlook theorist may not be perceptions proper, but rather perception-like representations with conceptual content.

The third chapter takes on the topic of emotion and considers the claim made by some philosophers that emotions are perceptions of evaluative properties.⁴¹ Although I ultimately reject this view, reflecting on emotions leads to consideration of constitutively phenomenal representational states. I speculate there may be evaluative properties most naturally represented by qualitative conscious modes of presentation. This idea ends up playing a role in the fifth and final chapter when I consider meaning in life and the question of what makes human life itself choiceworthy.

The fourth chapter looks at thought, understanding, and critical reason. The main tasks of this chapter are (i) to fit these higher level capacities into the broader picture of representational mind developed by the previous two chapters, and (ii) to explain how these capacities shape action in mature practical thought. This chapter develops a notion of intention as a conative state with conceptual content that functions to guide other conative states. This notion of intention informs to a discussion of rational motivation. It also reinforces my view of the overall structure of practical reason, which I develop more in the final chapter.

The fifth chapter tries to draw these threads together into a coherent (albeit rough) view of practical reason. I consider the notion of meaning in life, and develop the idea that meaningfulness is a value that functions to organize other values. I also try to integrate the most plausible account of meaningfulness with the notion of dialectical activity. The upshot is a

⁴¹ See Brady (2013) and Tappolet (2017) for discussion of this idea. I focus mostly on Tappolet's articulation of the view.

clearer view of the structure of practical reason and of the importance of particular projects, activities, and individuals (including other persons) in explanations of the choiceworthiness of human life itself. Practical reason has a hierarchical structure in which general goods, like living a good or meaningful life, are constituted by more particular goods, like having a career, having friendships, and so forth. Despite this structure, where the most general goods seem to explain the value of the more particular goods, I conclude that we cannot explain the value of the most general goods in life without an understanding of the particular goods that constitute them. This conclusion comes in large part from reflection on evaluative outlooks, especially the ineliminably phenomenal dimension of evaluative understanding needed to account for dialectical activity.

Chapter 2

Representation, Evaluation, and Agency

An evaluative outlook account of practical reason insists that rationalizing explanations of human action must appeal to ineliminably evaluative representational states. Moreover, many of these evaluative representations are products of relatively primitive subrational processes. This chapter lays the groundwork for defending an evaluative outlook model by developing the relevant notion of representation, presenting an account of evaluative content, and distinguishing between rational and subrational representational states. In pursuing these ends, I also begin to trace the relationship between representation and action.

The human mind has a variety of representational capacities: sense perception, emotion, and propositional thought, to name a few. These capacities involve distinct modes of presentation, and these modes of presentation help differentiate various types of representational state. Moreover, representation is closely tied to agency. From a phylogenetic standpoint, representational capacities enable more sophisticated and precise responses to an organism's environment—presumably that explains why they evolved. But this gives rise to a much-discussed philosophical puzzle. How exactly do representational states move an organism to act? In simple organisms, one can imagine algorithmic responses to different kinds of representation; for example, a representation as of danger could directly cause an organism to flee the object represented as dangerous. But in more complex organisms like humans (as well as other mammals, some birds, and perhaps a variety of creatures with more sophisticated psychologies) this kind of explanation seems to fall short. Humans can refrain from fleeing what

they know to be dangerous, refrain from eating what they know to be nutritious, and so forth. Can we expand our account of representational capacities to account for this (apparently) non-algorithmic agency?

Plausibly, the key capacity enabling this complex agency is conceptual representation—that is, *thought*. Thought, the ability to form propositionally structured representations (such as beliefs), which are incorporated into inferential networks, could create a layer of separation between lower level representation and action. Given the role of inferential networks in thought, if thought does indeed create the capacity for non-algorithmic responses to representations, then thought is the ground of *rational agency*.

I shall take as pretheoretical starting points both (a) that human beings, like the simpler organisms, have action-oriented, non-conceptual representations that occur (phylogenetically and ontogenetically) prior to the level of thought, and (b) that the representational capacity of thought is indeed non-algorithmic.⁴² Given the apparent connection between values and motivation, I suppose that these action-oriented representations are precursors to evaluations. Thus, I think these lower level representations belong in the broad category of *evaluative representations*.

From this initial sketch, a number of issues present themselves. First among them is to justify including perceptual level representations in the category of evaluative representation. Another relevant question concerns how, if at all, perceptions can have conceptual content. The interplay between more primitive representational capacities like perception and more sophisticated capacities like thought requires careful attention, especially if there are other representational systems that mediate between these two capacities.

⁴² For my purposes, non-algorithmic just means that thought processes cannot be modeled by a closed and complete computational system.

The key points for my case are that (1) concepts in humans are developed primarily through their application in propositional representations and inferences, which are constitutively rational processes, but (2) these concepts often are applied in context-bound reference through the same kind of (subrational) process through which higher-level perceptual attributives (“percepts”) are applied. The upshot for the evaluative outlook account is that there is a kind of evaluative representation poised to play a role in practical thought analogous to the role played by perception in theoretical thought. These perception-like evaluative representations provide the starting material for practical thought—the evaluative givens, so to speak. But rational processes, such as conscious reflection and inference, as well as experience can shape our understanding of evaluative properties, kinds, and relations. This evolving understanding in turn shapes the subrational processes that deliver perception-like evaluative representations, thus creating a kind of feedback loop that, when things go well, leads to improved evaluative understanding and improved practical thinking. This feedback loop is exactly what is needed to explain dialectical activity. I develop these ideas more in Chapter 3 and Chapter 4. For now, I focus on the nature of evaluative representations.

In the first section of this chapter, I discuss the psychological background. I define the notion of representation relevant to evaluative outlook models, and drawing from the work of Tyler Burge (2003, 2010a, 2014, 2022), I sketch a general account of perception as the most primitive kind of representation in nature. Then in the second section, I discuss evaluative representational content. I argue that a fundamental kind of evaluative representation serves the conative function of goal-setting within action commands, thus forging a link between evaluation and motivation. In the third section, I explore the possibility of representational systems that mediate between perception and conceptual thought. I argue that even though perception as such

cannot have conceptual content, there are almost certainly conceptual representations with perception-like structure, which share important features with perceptions within the space of conscious reasoning. I call these states perceptual cognitions. Moreover, there are important ways that capacities for thought shape perception, even if only indirectly. The upshot is that there are empirically viable ways that evaluative representations can be sufficiently perception-like to play the role required by the evaluative outlook view. Finally, I conclude by leveraging this account of evaluative representation to answer a general critique which accuses evaluative outlook theorists of being insufficiently precise about the kinds of apprehensions of value needed to explain rational actions.

I. Representational Mind

One might think that there is nothing special in the evaluative outlook model's insistence that rational explanations of action require appeal to evaluative representations. After all, traditional Humean models of rational action appeal to beliefs and desires, both of which are construed as representational. The difference between the evaluative outlook view and the Humean view lies in the kind of representational content that plays the central explanatory role. The Humean view banishes any ineliminably evaluative representational content. According to such a view, rational actions are explained by the presence of a *belief* that a course of action will bring about some state of affairs, and a *desire* that the state of affairs be brought about. The belief is a purely descriptive (non-evaluative) representation about the world: that such-and-such a state of affairs will be produced by such-and-such an action. The desire is another purely descriptive (non-evaluative) representation of the same state of affairs, along with some non-representational positive valence. Thus, the desire may have an evaluative element, but that evaluative element is

not part of the desire's representational content.⁴³ The evaluative outlook view rejects this picture, insisting instead that rationalizing explanations of action require appeal to representations with distinctively and irreducibly evaluative content. To unpack this claim, we must first develop an understanding of representation.

Representation is a distinctively psychological kind. At the most general level, representations are mental perspectives on an objective reality.⁴⁴ When they are successful, representations can both refer to and characterize entities in the real world. Tyler Burge (2010a: 292–308) argues that representation is a fundamental psychological kind because it figures ineliminably in the successful scientific explanations of perceptual psychology, and it cannot be reduced to other non-psychological notions, like biological fitness. The notion of representation is irreducible in this way because non-psychological notions cannot capture the relation of *veridicality*, or correctness, which is essential to the notion of representation as employed in perceptual psychology. Successful explanations in perceptual science depend on the notion of veridicality, a basic relation between an individual's mental states and its physical environment.

Perception is the most primitive capacity for objective representation by individuals (Burge 2010a: 367–436). Perceptions make reference to particular entities in an individual's physical environment and categorize them according to various attributes, such as *shape*, *color*, *size*, and *distance*. One of the constitutive features of perception is its immediacy: perception is always as of particulars in the individual's environment, and perceptual *attributives*—roughly, categories for sorting entities according to kinds, relations, and properties—always function to pick out particulars in the individual's environment. (There are complexities surrounding

⁴³ See Brewer (2009: 16–19) for exposition of this common view of desire and its role in action explanation. Brewer, of course, rejects this view.

⁴⁴ Non-mental conceptions of representation are, I think, parasitic on mental ones. For example, pictures or written words are representations only because they can be grasped by an individual with a mind.

perceptual memory and anticipation, but I set those aside for now. The basic point stands: perceptions are always context-bound.) To be veridical, the perception must be *caused* in the right sort of way by the particulars that the perception represents. If an individual perceives a small red sphere in its vicinity, that perception must be caused by the small red sphere and its features impinging in the right sort of way upon the individual's sensory systems for that perception to be veridical. Perceptions are individuated partly by their *veridicality conditions*: the conditions under which the perception is accurate. So a perception is veridical when the world is (approximately) as the perception presents it. And a perception is the kind of mental state that it is in virtue of having the specific veridicality conditions that it has. For example, a perception as of a red sphere is a different kind of mental state than a perception as of a blue cube.⁴⁵

Perceptions always apply attributives to particulars in order to pick those particulars out in context-bound singular (or plural) reference. Thus, a perception with the content red sphere there refers to the red sphere at the location (indexed relative to the individual) by applying those attributives (red, sphere, and the ego-indexed location). It can achieve this reference even if those attributives are non-veridically applied. Suppose for instance that through some optical illusion a blue cylinder appears to an individual to be a red sphere. The perceptual state still refers to that particular object by applying the (non-veridical) attributives as long as there really is an object causing the perception in something like the normal way. *Some* successful attribution is necessary for reference; in this case, there is a physical body at roughly the perceived location.

Perceptions proper, according to Burge (2010a, 2014), always occur in specific sensory modes: visual, tactile, auditory, and so forth. But perceptions are nearly always subject to

⁴⁵ The “as of” locution is meant to emphasize the non-factive nature of perceptual content. A perception as of a red sphere could be non-veridical—it could be that there is no red sphere there. Whereas a perception *of* a red sphere might imply the correctness of the representation.

intermodal influence (Burge 2022: 20–21). There are also *amodal* representations that occur at the level of perception—occurring alongside perceptual capacities, dependent upon perceptions for their representational content, and not dependent on more sophisticated cognitive representational capacities (Burge 2022: 21, 722–723). Pre-conceptual amodal representation takes inputs from particular perceptual systems (such as vision), but “abstracts from any specific perceptual modality” (Burge 2014: 2).⁴⁶ Modality is significant because perceptions have an *iconic noun-phrase-like* format, which means that they represent their subject matters through “natural correspondences” (non-representational relations of similarity) with that subject matter that enable direct reference to particulars—for instance, that x (Burge 2022: 293–363). To illustrate iconic format, consider a map, where particular features of the map correspond naturally to spatial relations of the territory being mapped. These are elements of perceptual modes of presentation.

Another more sophisticated level of representation involves propositional structures, where there are representational elements that do not primarily function to pick out particulars. For example, in propositionally structured representations, predicate-place attributives are applied to subject-place objects, but those attributions do not function to pick out the objects; reference is already achieved, and the attribution is separate.⁴⁷ For example, a thought with the content “The sphere is red” applies the conceptual attributive red through predication, and that attribution does not play a role in referring to the subject of the thought (the sphere). By contrast, in a perception as of a red sphere, the attributives *sphere* and *red* play an essential role in

⁴⁶ In some writings, Burge labels this level of representation “pre-conceptual cognitive” representation (2014: 2–3), but he later seems to question whether the label “cognitive” is appropriate for such representations (Burge 2022: 722–734). Regardless, he settles on a clear distinction between perception, understood as modal and intermodal sensory representation, and amodal representations, even when they occur at the perceptual level, on the grounds that amodal states are insufficiently “sensory” to count as perception proper (2022: 21).

⁴⁷ See Burge (2022: 202).

achieving the reference to the object. *Concepts* are representations that function constitutively to be elements in propositional structures, though they need not always occur within such structures (Burge 2022: 44). Propositions are representations with subject-predicate structure. Conceptual representational capacities draw from and depend upon perceptual representational capacities, although not all conceptual representations are based directly on perceptual representations. The reverse does not hold; perceptual capacities do not depend upon conceptual capacities, though they can be influenced by them in various ways.

Perceptual systems apply perceptual attributives algorithmically on the basis of proximal sensory stimulation, previous representational states, and what Burge (2010a: 376–396) calls *formulation principles* that function to overcome underdetermination.⁴⁸ The information registered from proximal sensory stimulation nearly always, even after accounting for previous representational states, underdetermines the veridicality conditions of possible perceptions that might be formed on the basis of those stimulations. Formulation principles allow the system to form perceptions that are, in the vast majority of cases, veridical perceptions despite this underdetermination. And perceptual illusions can also be successfully explained by appeal to these formulation rules. Importantly, on Burge’s view, formulation principles and the computations necessary for the formation of perceptual states occur at the *sub-individual level*. That is, they are *not* part of the representational contents of the perceptual states; nor need they be represented at all anywhere in the individual’s psychology. Therefore, even though perceptual

⁴⁸ Perceptual application of attributives is algorithmic in the sense that it can be mathematically described and successfully modeled computationally. On the other hand, there are no successful mathematical or computational models of conceptual representation or inference. This is not to say that such a model is impossible, but it would be a mistake, I think, to assume from the armchair that conceptual representation and inference can be computationally described.

states are (by and large) attributable to individuals, the processes through which perceptual states are formed are not; these processes are instead attributed to *subsystems* of the individual.

Burge (2010a, 2014) differentiates between perceptual systems and representational action systems. Representational action systems generate what he calls *actional representations* (2010a) or *conative states* (2022). These include goal-setting representations, which apply what Burge (2010a) calls *functional attributives*, and action commands, which map out and instigate particular bodily movements. Examples of functional attributives include things like danger, food, and mate. An example of an action command is a motor command to reach one's arm toward a food item. The veridicality conditions for action commands include the successful performance of the action specified in the representational content. Burge (2010a) describes the function of perception-level action commands as follows:

The simplest representational function of action guided by perception is to make veridical the actional representation that maps out the action. A frequently related representational function is to fulfill through an action a perceptually set goal. The action is to fit the representational content that maps out the action, or the goal. (339–340).

Goal-setting conative states have at least two components. First, applications of functional attributives, like food, function to pick out objects in the environment that are apt to be the targets of action commands. Second, imperatival action schemas specify generic action types with respect to those targets, such as to eat that food.

Burge (2010a, 2014, 2022) treats perception, perceptual memory, perceptual anticipation, and pre-conceptual conative states as all being at the same representational level, below cognitive representations. In his view, these perceptual-level representational states all lie at the foundation of what it is to have a psychology, and they all occur together in an animal's

psychology if any occur at all. Nonetheless, Burge distinguishes between *lower-* and *higher-level* attributives within modal perceptual systems.⁴⁹ He focuses on vision. Shape, color, and motion are examples of lower-level attributives that occur in the visual system. Higher-level visual attributives are formed on the basis of, or *through*, lower-level visual attributives:

All causal factors in normal formation of higher-level perceptual attributives go through low-level perceptual attributions. Suppose that body, face, and dangerous are higher-level visual-perceptual attributives. A visual system cannot produce a representation of something as a body, or as a face, or as dangerous without attributing a shape or motion by which it is perceived as a body, a face, or dangerous. (2014: 3)

Moreover, low-level perceptual attributives can form the basis for other high-level attributives even outside the perceptual system:

On the basis of vision alone, one can represent something as a pushing, as an instance of agency, as a pine tree, as a piano, or as one's favorite movie villain—*only* by visually attributing low-level attributes. Some higher-level attributives for the listed attributes are probably produced in the visual system as well as in conceptual systems. Others are produced only posterior to perception.... (3)

This latter point about the perceptual basis of higher-level attributions is important, and I will return to it. But for now note the implications of these points for the question of whether conative states are more primitive than non-conative perceptions, as some thinkers have proposed.⁵⁰ If Burge is right, then conative states are *not* more primitive than non-conative

⁴⁹ See Burge (2022: 118–120, 124–130).

⁵⁰ See, for example, Campbell (2011). Campbell claims that “affordances”—what Burge (2010a) calls functional attributives—are all that primitive organisms perceive, and he presents this claim as a direct rejection of Burge’s view that such attributions require more basic attributions of physical properties, like shape and motion. The nub of disagreement seems to be that Campbell does not think those lower-level attributions, upon which the higher-level attribution of affordances depend, are represented by the individual, but only by the individual’s subsystems. Thus, the disagreement between Campbell and Burge seems to turn on a rather fine point. Burge thinks that empirical explanations in vision science present lower-level attributions as being made at the whole-individual level. See also Noë (2004) for an account of perception that places emphasis on perceiving affordances for action rather than lower-level physical properties. Noë, however, is less clear about whether he thinks primitive organisms perceive *only* affordances.

perceptions. If functional attributives like danger are always formed on the basis of lower-level attributions like shape and size, then non-conative perceptual capacities are necessarily operative wherever capacities for functional attributions are—even though functional attributional capacities and conative systems may well have evolved in tandem with each other.⁵¹

We can, of course, imagine an organism that forms representations as of danger, for example, as immediate responses to sensory information registration. Such an organism would form representations as of danger without also attributing shape, size, motion, etc. But I doubt that sensory information registration would directly lead to applications of functional attributives without also being used in the application of non-functional attributives, like shape or distance. The reason for this doubt is that it is unclear in such a case what role the representational system plays that could not be accounted for in non-representational terms. If an organism senses light, say, and reliably moves away from electro-magnetic radiation of certain frequencies, that behavior might be explained without invoking any representational capacities at all. If, however, the organism moves away from the source of the light in a way that is sensitive to the exact location and patterns of movement of that light source, that behavior suggests that the organism has representational capacities since the behavior is best explained by invoking perceptual constancies and perceptual anticipation. For example, if the organism not only moves away from the light source, but does so more quickly when the light source is moving, and in a way that anticipates where the light (likely) *will be*, these behaviors cannot be explained merely in terms of sensory registration of occurrent electro-magnetic frequencies and intensities.

The point here is that the evidence that suggests the presence of a representational system involves responsiveness to particulars in the environment that cannot be accounted for without

⁵¹ Burge makes an argument to this effect in (2022: 118–124).

positing representational attributives other than the functional ones—in addition to a representation as of danger, the organism would also need to be able to represent things like movement and distance. Another way of putting the point is that conative states could not serve their function without non-conative representational attributives, like size and distance.

It is important to keep in mind the distinction between goal-setting representations (like perceptions that apply functional attributives) and commands.⁵² The veridicality conditions of representations that apply perceptual functional attributives include teleological relations between the individual and its environment. For example, a perceptual-level representation with the content danger there is veridical when the represented entity does in fact pose a threat to the well-being of the organism, and thus if it would in fact be *good for* the organism to flee (or otherwise avoid) the represented entity (other things being equal). Thus, I think there is a sense in which such perceptual representations may be considered *evaluative*. I do not mean that the

⁵² Burge thinks it is an open question whether functional attributives like danger occur in perceptual systems or only in actional systems (2010a, p. 280 fn. 142). Presumably, an actional system takes inputs from the organism's various representational systems (perception, belief, and so on) and generates conative states in response to those inputs. Suppose that there is a direct causal path from a representation as of danger, for example, and an organism's action commands to flee. The variables that determine whether the organism attempts to flee are narrowly definable. Whenever the attributive danger is applied by the organism's representational systems—and perhaps some other narrowly specifiable conditions are met, such as the seemingly-dangerous entity being represented as moving towards the organism, for instance—the organism will form the appropriate actional representation and attempt to flee. The fleeing can be understood as a behavioral protocol instigated by representations starting with the functional attributive. It may be that in simple organisms the perceptual system and the conative system are not separate in any meaningful sense. In more complex organisms, a separation between conative systems and other representational systems seems more likely. For one thing, when there are multiple representational systems that bear on action (perception and thought, for instance), a separate conative system that takes inputs from the various non-conative systems would neatly explain, at least in the abstract, how those separate representational systems could coordinate to produce behavior. Moreover, the apparent ability of systems of thought to intervene between applications of perceptual-level functional attributives and actions—as when a person refrains from fleeing something she perceives as dangerous—suggests a separate system that takes inputs from the other representation systems. That thoughts sometimes succeed and sometimes fail to guide action in cases where thought and perception conflict further supports such a conclusion. Moreover, it seems possible that evaluative attributives like danger occur only amodally. Since action-oriented representations function to guide the behavior of the whole organism, it may be that they always abstract from any particular sense modality. At any rate, I think that *some* evaluative representations almost certainly do occur at the pre-conceptual cognitive level. For my purposes, I am not sure how much really rides on whether attributives like danger occur within the visual system, for example, or only at the levels of pre-conceptual conative representations, pre-conceptual cognition, and propositional thought.

representations have any notion of good or value as part of their content. They don't.

Nonetheless, their veridicality conditions cannot be specified without reference to the good of the individual with the representation.⁵³ Thus, I sometimes call such functional attributives “proto-evaluative” to indicate their being at the perceptual level and their having no explicit evaluative content.

Commands, on the other hand, have a less direct connection to value. A command maps out a course of action for the organism, and it must be *made* veridical by the organism's movements. For this reason, I say that commands are future-directed. Their veridicality conditions are specified in terms of the organism's behavior *after* the representation is formed. Moreover, the veridicality conditions of commands come apart from the good of the individual in a way that veridicality conditions of low-level evaluative representations do not. A particular command is functionally evaluated according to how well the organism's behavior matches the action mapped out by the representation, regardless of whether or not that action serves the good of the organism. For example, an organism might, because of disease or malfunction, form a command to destroy part of its own body. Presuming that the bodily destruction does not serve some further functional purpose (like survival), the action, if performed, would be entirely to the detriment of the individual's flourishing. Yet nevertheless, the command *itself* is made veridical by the execution of the act of bodily destruction that it maps out, and it would be a kind of *actional* and *representational* failure if the organism for some reason did not perform the act of bodily destruction so mapped (perhaps because of motor paralysis, for instance). It would be such a representational failure even though it is good for the organism in an overall sense that it

⁵³ Unless otherwise noted, I use the term “good” here in the sense of well-being or functional success. I use the term “worth” to indicate value that is not relative to function or well-being, as when some entity is valuable for its own sake, an end in itself. I take the term “value” to be generic, encompassing both of these notions.

not perform the action—that is, it would be a success in terms of individual well-being or flourishing.

This distinction between evaluative (goal-setting) representations and commands goes beyond the level of perception. As I use the terms, *desires* (conceptual representations that function to be actional goals) are evaluative representations, while *intentions* (conceptual representations mapping out particular courses of action) are commands. Both these kinds of representation—desires and perceptual goals on the one hand, and intentions and perceptual commands on the other—are *committal* in the same way that beliefs are committal while imaginings are not, and that willings are committal while wantings are not. In general, committal representations constitutively function to answer to reality.⁵⁴ For example, beliefs are committal because they constitutively function to be true, and a belief fails *qua belief* if the world is not as the belief presents it. Similarly, a desire is committal because it constitutively functions to be a correct evaluative representation, and it fails *qua desire* if the object of the desire does not have the value the desire presents it as having. There is a different sense of “committal” in which desires and intentions differ. Intentions (and other commands) are committal in a future-directed way; an intention that is not carried out fails *qua intention*. Evaluative representations, like desires, do not “commit” the individual to any course of action in the way that commands do.

Evaluative representations, at least in the human case, clearly do not entail (either conceptually or causally) commitment to any action at all. The same may be true for organisms with merely perceptual psychologies. Suppose an evaluative perception, with content like danger there, need not by itself prompt the formation of any command. Thus, even under the assumption

⁵⁴ See Burge (2022: 191) and (2010a: 74–75).

that perceptual systems are entirely algorithmic, the link between applying a functional attributive like danger and taking the appropriate action may not, or at least not always, be direct.

Functional attributives (proto-evaluative representations) constitutively function to serve as goals (or targets) for action commands. But commands also include representations of action types as well as goals (Burge 2022: 508–509). For example, objects flagged by the perceptual system as dangerous typically become the objects of commands to flee. Such a command might have a content like flee that danger. Such a command, in conjunction with other perceptual-level representations (including occurrent perceptions, short-term perceptual memories, and perceptual anticipations of both the target and of the individual’s own bodily positioning), causes and guides lower-level motor commands that direct specific bodily movements.⁵⁵

Functional attributives are tightly connected to action categories. For example, danger and to flee are connected. It may even be that the functional attributive danger is functionally equivalent to to be fled. That is, the representational content of a perceptual representation as of danger is partly constituted by the functional role played by that attributive in the relevant commands. Nevertheless, the veridicality conditions of functional attributives in perception cannot be specified entirely in terms of relevant actions; they must be characterized in terms of the good of the individual. Whether or not an object is veridically categorized as to be fled, for example, depends on whether that object in fact poses a threat to the individual.⁵⁶

Useful applications of such functional attributives are not always veridical. It sometimes serves the organism’s biological interests to (have a tendency to) apply functional attributives non-veridically. For example, consider the case of a rabbit that more or less indiscriminately

⁵⁵ See Burge (2022: 512–513).

⁵⁶ I think this parallels the way in which desires cannot be specified entirely in terms of reasons for action. See Brewer (2009: 153–157).

perceives shadows cast from objects passing overhead as dangerous. Let us suppose that *any* overhead shadow moving on the ground results in the rabbit forming a representation as of danger, and moreover that the rabbit *always* attempts to flee such shadows.⁵⁷ It is not hard to imagine how such over-sensitivity in the rabbit's perceptual system and conative system could serve the function of biological fitness; because of the speed at which birds of prey can move, rabbits that attempt to flee at the slightest sign of such a predator may be more likely to survive and reproduce. But when the rabbit forms a representation as of danger applied to a cloud passing overhead, that representation is non-veridical; the cloud is not dangerous to the rabbit. And thus there is a sense in which the rabbit is making a mistake when it attempts to flee the cloud.

In contrast, imagine a rabbit that applies the attributive danger in just the same way, but that does not immediately attempt to flee in response to those attributions. Perhaps there are additional conditions that must be met before the rabbit attempts to flee. For instance, perhaps the rabbit pauses, fixes its perceptual attention on the object perceived as dangerous, and then flees if, say, certain additional more fine-grained movement patterns are perceived, or some more specific size and shape is perceived, and so forth. (I am assuming that all of this plays out in the rabbit's perceptual systems; in these examples, the rabbit's capacity for thought, if it has such a capacity, does not play a role.) This rabbit, unlike the first, may be more likely to avoid the error of fleeing things that are not actually dangerous, but perhaps less likely to successfully escape the real threats—fast-moving birds of prey.⁵⁸

⁵⁷ I am not claiming that rabbit perceptual systems actually function this way; the example is hypothetical for the purpose of illustration. But it is likely that many organisms, including rabbits, have such unreliable perceptual capacities that can be explained by various evolutionary pressures (see Burge 2003: 517).

⁵⁸ As Burge (2022: 507) puts the point, “If danger looms, fast action in response to generic type is better than slow action in response to detailed categorization.” As Burge discusses in detail, perceptual processing often makes generic categorizations *before* more specific ones.

Of course, we can see why evolution might select for the former system, which sacrifices some representational success for biological survival, rather than the latter system; however, the example illustrates the important difference between representational and biological function, and the complexity that arises because of the functional connection between representation and agency. Representational functions evolved to serve biological functions, but they are not identical to them.

The upshot here is twofold. First, a primary *representational* function of attributives like danger is to be veridical, which means correctly indicating objects that in fact pose a threat to the individual organism. In this way they are exactly like other non-functional attributives. Another function of these proto-evaluative representations is to serve as goals for action commands. This is separate from their representational function to be veridical. A functional attributive can serve as a goal for a command even if it is non-veridically applied, as in the case of the rabbit fleeing the cloud, which it incorrectly perceives as dangerous.

Second, there is apparent continuity between the perceptual and conceptual levels of evaluative representations and actional representations. At both levels, evaluative or proto-evaluative representations (applications of functional attributives at the perceptual level, desires at the conceptual level) function to be veridical *and* to serve as goals in action commands. At both levels, commands can be formed and enacted (that is, made veridical through action) even if the evaluative representation that serves as the goal is non-veridical. Just as the rabbit can see something benign as dangerous, a human being can desire something that is bad, or that lacks the value the desire presents it as having. According to evaluative outlook accounts, both of these representations—perceptual and conceptual—are explained as instances of something *seeming* good or bad to the individual. At the perceptual level, of course, that

evaluative seeming does not have any explicit evaluative content; it is only evaluative in the sense that the functional attributive's veridicality conditions include aspects of the individual's well-being, flourishing, or well-functioning. At the conceptual level, a desire may or may not have explicitly evaluative content.

Returning to our rabbit, the case suggests at least three kinds of possible error to consider: (1) representational error, in which an attributive is non-veridically applied; (2) actional error, in which the enacted bodily movements do not match the action mapped by the relevant action command; and (3) functional error, in which the individual forms and enacts a command aimed at a non-veridically represented functional goal.⁵⁹ The first two of these functional failures are independent of biological function; representational functions such as these do generally serve biological functions of survival, flourishing, and fitness for reproduction, but they are not the same as those biological functions and can come apart from them in many cases.

The third kind of error—functional error—arises when representational error occurs alongside actional success. The nature of this kind of error depends on the nature of the evaluative attributive applied (non-veridically). For example, perceptual functional attributives always have *some* kind of biological well-functioning as an element of their veridicality conditions. Perceptual attributions of danger are veridical when they are applied to objects that in fact pose a threat to the individual's survival, flourishing, or fitness for reproduction. So when our rabbit attempts to flee a cloud, an object that it incorrectly perceives as dangerous, the

⁵⁹ It is an interesting question whether the mere formation of an action command with a non-veridical functional attributive as the goal is a kind of error. Does the rabbit, when it forms the command to flee after perceiving the cloud (or the cloud's shadow) as dangerous, make another mistake, in addition to the non-veridical application of the functional attributive, and separate from the actional failure of fleeing something not dangerous? In other words, is it a *representational* failure for an action command to be formed as the result of a non-veridically applied functional attributive? Given that functional attributives like danger function to serve as goals in commands, it is not clear how to answer this question. Fortunately, I do not think anything terribly important hangs on it.

representational failure engenders a biological failure as well, since the biological function of the action is to escape objects that are in fact dangerous. The *particular* action is a biological failure in this regard even though the *pattern* of action—in which the rabbit often attempts to flee even when there is no threat—is a success in terms of increasing the likelihood of survival and reproduction.

And of course, even within purely biological notions of well-functioning, these functions come apart. It may, for example, serve the rabbit's biological goals of survival and reproduction to have a representation system prone to this kind of error, but it also works against the rabbit's biological flourishing, since the rabbit presumably would be better off if it spent less energy fleeing, and only attempted to flee actual threats. Obviously, the rabbit must survive in order to flourish. But it would be *better* for the rabbit, in terms of flourishing, if it had a more discriminating representation system that could better distinguish threats from non-threats and that still allowed the rabbit to flee real threats quickly.

More sophisticated representational systems, like thought, enable evaluative attributions connected to other notions of value besides biological success. I discuss conceptual evaluations more in later chapters, especially Chapter 4. For now, I turn to a more detailed examination of the content of evaluative representations.

II. Evaluative Representational Content

Representational content is constituted by (i) the specific veridicality conditions of the representation, and (ii) the mode of presentation, or the way those veridicality conditions are presented in the mind of the individual. Thus, two different perceptual states can have the same veridicality conditions, but different content because they have different modes of presentation.

For example, a perceiver can represent the shape and location of an object both tactilely and visually. The tactile representation and the visual representation will both have the same veridicality conditions—both perceptions, since they are as of the same object, are veridical only if there is in fact an object of such-and-such a shape at such-and-such a location, and that object is the causal source of the representational state in the right sort of way. But the two perceptual states will have different modes of presentation.

The focus of most mainstream philosophical discussions of perceptual content is perceptual *experience*.⁶⁰ I think that this emphasis is unfortunate, insofar as it obscures the distinction between representation and consciousness. As a pre-theoretical starting point, the two clearly come apart. It is conceivable that there are organisms capable of representation but that do not have conscious experiences; it is also conceivable that there are organisms without representational capacities that do have conscious mental states (though this claim is harder to assess since we know so little about the constitutive conditions of consciousness). Moreover, there are examples of human mental states—such as pain and pleasure states—the natures of which are constituted by conscious qualitative elements and not by representational content. Tokens of these states may have representational content, but such a token is not the kind of state that it is in virtue of having that content; rather, it is the kind of state that it is in virtue of its phenomenal qualities. For example, a mental state is a pain state if and only if it has certain phenomenal qualities (pain qualities). Perhaps some pain states also have representational content—of tissue damage, for instance. But that representational content is neither a necessary nor a sufficient condition of the mental state being an instance of pain. A mental state with no representational content whatsoever, but that has the phenomenal qualities characteristic of pain,

⁶⁰ See, for example, Crane and Finch (2021) and Siegel (2021).

is a pain state; and a mental state with representational content as of tissue damage, but that lacks the phenomenal qualities characteristic of pain is not a pain state.⁶¹

The reverse is also true. There are human mental states that, although often phenomenally conscious, have natures that are constitutively determined by their representational content, and not by their phenomenal qualities. Perceptions and beliefs are examples of such states. A visual perception as of a red sphere has a characteristic phenomenal quality (at least in humans, and probably in other similar organisms). But it is a perception as of a red sphere if and only if it has the representational content red sphere. The characteristic phenomenology of seeing red, for example, is not a necessary condition of its being a perception as of redness. Moreover, there is clear empirical evidence of human representational states that occur in the absence of reflectively accessible consciousness (Burge 2010a: 188). Of course, it may turn out that representational capacities are a necessary precondition of phenomenal consciousness, or that phenomenal consciousness arises with all representational capacities. The relationship between representational capacities and phenomenal consciousness is unclear. But these are empirical matters and cannot be settled from the armchair.

It is true that phenomenal qualities are often elements of perceptual states in humans (and probably in many other animals). These phenomenal qualities may even be characteristic modes of presentation for many perceptual states. For example, the qualitative experience of seeing red may be the characteristic mode of presentation of visual perceptual states as of redness in humans and other mammals. But as a constitutive matter—that is, a question of the necessary or

⁶¹ One might think that some pain states are not always felt. For instance, I might have a pain in my leg that I do not immediately recognize as a pain in my leg. I could realize in the middle of the afternoon that my leg has been in pain all day, even though I was not aware of the pain until this moment. I do not think that this kind of case undermines the point here. Once I am aware of the pain, it has the phenomenal quality characteristic of pain, and it is in virtue of this quality that the pain is a pain. This is true even if lack of attention or other mechanisms make it so that I do not feel the pain at every moment that I am in the pain state.

sufficient conditions for mental states being the kinds that they are, and having the natures that they have—there are mental states, such as perceptions and beliefs, that are constitutively representational. The nature of such a state is constituted by its representational contents, and not by its phenomenal presentation, if it has any. Similarly, there are mental states, such as pain states, that are constitutively phenomenal. The nature of such a state is constituted by its phenomenal qualities, and not by its representational contents, if it has any. It may be that in beings with phenomenal consciousness many evaluative representations have characteristic phenomenal qualities, similar to the way that color perceptions do. But in keeping with my working hypothesis that some organisms have perceptual capacities but lack phenomenal consciousness, I will not assume that the modes of presentation of all evaluative representations are constitutively phenomenal.

I have already mentioned the veridicality conditions of one kind of evaluative representational content. Functional attributives, like danger or food, have veridicality conditions specified by reference to the well-functioning of the organism. Presumably these (proto-)evaluative perceptual states also have characteristic modes of presentation. For example, a perception of danger presents its object in a *dangerous way*, just as a perception of a sphere presents its object in a spherical way. If it turns out that functional attributives occur within particular modal perceptual systems, like the visual system, then there may be different modal ways of presenting danger, food, and so on. It may be, however, that evaluative perceptions are always (or primarily) intermodal. Perceptual-level conative representation systems are probably intermodal; if functional attributives only occur in conative systems, then they will not occur in specific sensory modes.⁶²

⁶² If Burge (2022: 513) is correct that conative states are always, or nearly always, intermodal, then once a functional attributive becomes part of a conative state, it must be abstracted from any one sense modality.

So far, I have mostly discussed perceptions that apply *functional attributives* within perceptual systems (or perhaps in perceptual-level conative systems): attributives like danger and food. These attributives function constitutively to be goals within perceptual-level action commands. I have called representations with this kind of content *proto-evaluative*, since at the perceptual level representations cannot have explicit evaluative content. On the view under consideration, functional (proto-evaluative) attributives are formed on the basis of lower-level attributives for properties like shape, size, distance, and so forth. These functional attributives are evaluative in the sense that their veridicality conditions depend on facts about the well-functioning or well-being of the individual. This means that we cannot assess the veridicality of a representational state that applies such a functional attributive without some understanding of the well-functioning of the individual whose representational state it is. I call this kind of evaluative content *implicit evaluative content*. The individuals whose representational states include such implicitly evaluative attributives need not be able to represent their own well-functioning. They need not be capable on any level of representing goodness, value, or worth as such.

There are other kinds of evaluative content, however, involving conceptual evaluative attributives, and sometimes explicit notions of value. Evaluative representational content containing explicit notions of goodness or worth I shall call *explicit evaluative content*. These representations are always conceptual, so no perceptual-level representations have explicit evaluative content. For the time being, this distinction between implicit and explicit evaluative content can be pushed to the background. I highlight the distinction here in order to avoid overintellectualizing perceptual capacities. Organisms with perceptual capacities can have

evaluative representations, but that does not entail that a notion of value is part of the content of any of their representational states.

But what about the possibility that perceptions can have conceptual content? If this is possible, then there may be more to perceptual evaluative representations than the application of functional attributives by perceptual systems or perception-level conative systems. We should, therefore, consider whether and how conceptual evaluative content could be perceptual.

III. The Space Between Perception and Thought

In individuals capable of conceptual thought, some representational states apply evaluative *concepts* on the basis of perceptual attributives. It seems upon initial reflection that we can perceive something as not just edible, but delicious, for example. Similarly, we can perceive an object not just as a round red object, but as an apple. Perceiving something as an apple, or even as a delicious apple, often appears introspectively to be a direct application of a conceptual representational category, *not* a separate judgment on the basis of more basic perceptions or perceptual beliefs. Whatever capacities apply the conceptual attributives appear to do so involuntarily and unconsciously. Introspection, however, is of limited use in delimiting psychological kinds, so we will need to reflect on our understanding of perception, thought, and conation.

There are at least three ways to explain such apparently cognitively-laden, perception-like representational states. First, it could be a case of genuine cognitive penetration. That is, it could be a case where attributives from a more sophisticated representational capacity—conceptual thought, or judgment—are actually applied *by the perceptual subsystem*. A second possibility is that these are really cases of perceptual belief in disguise. That is, the system of conceptual

thought applies propositionally structured representations on the basis of non-propositional (and non-conceptual) perceptual representations, and this process occurs quickly, automatically, and in a way that is not consciously or reflectively accessible to the individual. A third possibility is that some intermediate representational system, which operates posterior to perception but prior to conscious thought, applies conceptual attributives (attributives developed as part of a system of propositional thought) directly on the basis of lower-level perceptual representations. Depending on the nature of this intermediate representational system, we might think of this as a qualified kind of cognitive penetration.

The first and third possibility rely on the idea that conceptual attributives can be applied in representations that do not themselves have propositional structure. The second possibility, on the other hand, need not imply that conceptual attributives occur anywhere other than in propositionally structured representations. It is also worth noting that these possibilities are not mutually exclusive. It could be, for example, both that an intermediate representational system (at least sometimes) applies conceptual attributives, *and* that perceptual beliefs (at least sometimes) are formed quickly, automatically, and with reflective inaccessibility.

Explaining cases of apparent cognitive penetration of perception is an empirical matter. It certainly cannot be settled by introspection alone, and I cannot settle it here. But the correct explanation may have some implications for ethics. Explaining conceptual evaluative representations will be a crucial component of any successful evaluative outlook account of practical reason. If it turns out that all conceptual representations are propositionally structured thoughts, that would potentially undermine a key claim of the evaluative outlook view: that some mental states are perception-like representations of value, which can be shaped over time by experience and reflection, and which can serve as basic inputs for practical thought. But I think a

defeasible case can be made that some non-propositional and relatively low-level representational states apply conceptual attributives, and that these states, whatever their exact nature, can play the role required by an evaluative outlook account.

By *propositional representation*, I mean a representational state with at least two properties. First, propositional representations have conceptual content. That is, they apply conceptual attributives either to particulars, or to contextually uninstantiated kinds, properties, or relations. Attributives, whether conceptual or perceptual, always indicate kinds, properties, or relations. But perceptual attributives always and only apply to particulars as instantiations of the kind, property, or relation that the attributive indicates. This is what is meant by the claim that perceptual representations are *context-bound*. Conceptual attributives, on the other hand, may or may not be context-bound. Conceptual thought can refer, for example, to a specific red apple, or to redness, apples, or red apples in general.

The second essential aspect of propositional representation is *propositional structure*. Propositional structure, as I understand it, is subject-predicate structure. Subject-predicate structure involves a noun-phrase-like element and a main-verb-phrase-like element (Burge 2022: 44, 49). This structure creates a separation between the subject-place entity and the property, kind, or relation attributed to it. Unlike in perception, where all attributions occur as part of direct reference to particulars, subject-predicate structure allows reference and attribution to come apart. This separation creates the space for inference, since the elements within the structure are interchangeable with other elements. (Not to mention that predicate structure can instantiate logical relations.) For example, the thought *that apple is red* allows for transformations to *that fruit is red* or *at least one fruit is red*, and so forth. I think that both the subject-predicate structure *and* the conceptual contents within that structure are necessary for

inference. The multi-part structure allows for interchangeability of elements, and the conceptual content allows for generalization and logical transformations.

The hallmark of conceptual attributives is that they function constitutively to be constituents in propositionally structured representations (Burge 2010a, 2010c, 2022). But this constitutive point does not, I think, imply that conceptual attributives always and only occur within propositional structures. There are some representations with conceptual content but without subject-predicate structure. For example, I can think about elephants without predicating anything of them. More specifically, I think that some mental states apply conceptual attributives in perception-like structures. Perceptual structure is noun-phrase-like, and always includes an indexical element (Burge 2022: 49–50). Perceptual structure lacks a main-verb-phrase-like element, however. So a conceptual representation with perception-like structure could have a form like elephant there or that elephant. Moreover, many conceptual attributions directly import the noun-phrase-like iconic mode of presentation from perceptual representations. Often, one can think about that elephant because one perceives that elephant. This combination of perceptual form with conceptual content could help explain how perceptual-level conative systems and cognition work together to produce actions in organisms like human beings. But it could also help explain how human action is both similar to and different from action in more primitive organisms. It could explain why many of our evaluative representations are formed through processes that are reflectively inaccessible, but are nevertheless sensitive to our understanding, which *is* accessible to reflection.

I offer two general reasons to support the idea that perceptual structure can carry conceptual content. First, conceptual attributives are capable of context-bound reference fixing, just as perceptual attributives are (Burge 2022: 37). And context-bound singular reference is the

basic form of perceptual representation. Second, in humans (and probably many other animals), the relationship between higher order cognition and perception is complex, and there is much we do not understand. In the absence of conclusive evidence to the contrary, we should be open to possibilities for representational systems at different levels to interact. In the remainder of this section, I briefly expand on these points, and conclude that insofar as some evaluative representations have conceptual evaluative content and perception-like structure, such representations can satisfy the evaluative outlook model's requirement that there are subrational evaluative representations with ineliminably evaluative content. I call these subrational evaluative representations *perceptual cognitions*.

A primary structural difference between perceptual representations and propositional ones is that perceptual representations are always context-bound, and they always apply attributives to particulars partly as a means of picking those particulars out.⁶³ Thus, perceptual attributives always function, in part, to achieve perceptual reference. For example, in a perceptual representation with the content red sphere there, the attributives red and sphere are applied in order to pick out the object. The object is referred to *through* the application of the attributives. In propositional representation, attributives need not function to pick out particulars in this way. For example, in the thought content the sphere is red, the conceptual attributive *is red* does not function in its application to pick out the object. But *some* conceptual attributives in propositional representations do directly pick out particulars. In the same thought, the concept *sphere* could function in this way; the object is referred to *through* the application of that conceptual attributive.

⁶³ See Burge (2022: 44).

Importantly, this process is not the same as reference fixing through description. If it were, the reference would fail if the attribution were non-veridical, but that is not the case (Donnellan 1966). The thought content the sphere is red can succeed in referring to its subject even if that subject is not really a sphere, and is not really red. Similarly, a perceptual representation as of a red sphere (with the content red sphere there) can succeed in referring to its object even if that object is not really a red sphere. As noted above, *some* veridical attribution is necessary for reference to succeed, but many of the attributives applied by the representational state can be non-veridical and yet the state can still succeed in referring.

Since propositional representations include direct attributions of conceptual attributives that function to pick out particulars, it seems that propositional representations have elements with the same kind of structure as perceptual representations. Indeed, the most primitive kind of conceptual thought imports representational content, including modes of presentation, directly from perception (Burge 2010a: 544–547, 2022: 202). And if it is possible for perception-like structure to carry conceptual content *within* propositional structure, then it seems likely that there are perceptually structured representational states with conceptual content *outside* of propositionally structured states.

The second point I want to make touches on the debate over cognitive penetration of perception, or what are sometimes called “top-down” effects. The literature on this topic is vast and the issues at stake delicate.⁶⁴ I follow Burge (2010a, 2014, 2022) in defining *perception* as a psychological capacity for *sensory* representation. This understanding of perception rules out certain kinds of cognitive penetration of perception. In short, since perceptual capacities are

⁶⁴ See Tumulty (2020, Chapter 3) for a discussion of the empirical and philosophical literature on cognitive penetration and how the issue bears on ethics and epistemology. For additional critical discussion of that literature, see Burge (2022: 433–446, 677–680).

constitutively sensory, perception can only apply attributives for sorting environmental particulars according to sensorily discernable properties: size, shape, color, motion, and so forth. Perceptual systems, therefore, cannot apply attributives for properties such as internal mental states, which cannot be directly discerned by the senses. And higher-level attributives—such as body, face, and danger—are known to occur within perceptual systems precisely because the perceptual systems are capable of distinguishing these attributes from other “look-alike” attributes that co-occur in the environment.⁶⁵

Our understanding of what can and cannot be a perceptual attributive may be incomplete, but it is highly unlikely that human perceptual systems contain attributives for a sufficiently rich array of evaluative properties, kinds, and relations to do the work required by the evaluative outlook view. The subrational evaluative representations that the evaluative outlook model of practical reason insists on must include or be deeply connected to conceptual representations. Otherwise, the kinds of explanations of rational action that the evaluative outlook model offers simply would not be convincing. I think that the key point—what makes these representations subrational—is that the relevant evaluative representations are formed quickly, non-inferentially, and through automatic unconscious processes that are not necessarily attributable to the individual. It is in this sense that evaluative outlook theorists are justified in calling such representations perception-like.

A good candidate for this kind of perception-like evaluative representation is perceptually enabled conceptual sorting. This is where an individual applies a conceptual attributive directly on the basis of lower-level perceptual attributes, such as size, shape, and motion. Importantly, this capacity is heavily influenced by expertise and experience, even though it does not involve

⁶⁵ See Burge (2022: 120–124).

actual penetration of conceptual attributives into perceptual systems (Burge 2022: 436). For example, an experienced bird-watcher can identify different varieties of birds quickly and without conscious effort, based just on visual perception. The bird-watcher's visual system probably does not have attributives that pick out things like robins and finches as such.⁶⁶ But the visual system can be trained to sort objects in the visual field according to complex combinations of low-level perceptible properties (size, shape, color, and motion, for instance). So the bird-watcher's visual system may very well acquire an attributive that functions to pick out finch-looking objects.⁶⁷ The difference between this perceptual attributive and the conceptual attributive finch is that the veridicality conditions of the perceptual attributive are satisfied whenever the object in fact has the relevant low-level perceptible properties. The conceptual attributive is only veridical if its object is in fact a finch. Once again, perceptual learning can hone the individual's sensitivity to differences in perceptible properties to a fine degree, so an individual can become very good at perceptually picking out things like finches. But perception, because it is sensory in nature, cannot be sensitive to a thing's underlying nature in the way that conceptual thought can be. I will return to this point later, especially in Chapter 4.

The perceptual learning that enables this kind of conceptual sorting is certainly influenced by higher-level thought, at least indirectly. An individual with a fine-tuned perceptual ability to pick out finch-looking objects, for example, may acquire that ability because of

⁶⁶ Perception is unlikely to pick out natural biological kinds as such because the nature of a biological kind depends on features—genetic structure, evolutionary history, and so forth—that are not perceptible in a sensory mode. Thus, perception does not have the tools to be sensitive to differences between members of a biological kind and non-members that nonetheless have similar perceptible properties (size, shape, color, motion, and so on). But that sensitivity is the key empirical question, and it is possible that perceptual systems could acquire it, at least in some cases. See Burge (2022: 118–124).

⁶⁷ See Burge (2022: 427–431, 625–630).

patterns of attention that are caused by thoughts and intentions about finches.⁶⁸ Thus, even if cognitive content does not directly influence perceptual processing, there are ways that perception can nonetheless be indirectly structured by conceptual thought.

Consider what Brewer says about the conceptual structuring of perception and desire:

It is a perfectly general truth that our immediate apprehension of the world is structured by the concepts we have mastered, and it continues to take on a more refined structure as we master a more complex and ample array of concepts. The untrained eye looking at neurons through an electron microscope will see undifferentiated blobs with no more intelligible form than splattered paint. A trained neurologist will see synapses, axons, dendrites, and growth cones. This conceptual structuring does not seem to emerge as the product of judgments to the effect that certain blotches of color are synapses, axons, dendrites, or growth cones—they are seen as such. Indeed, it would be an unstable position to insist that concepts can only structure our experience via implicit judgments to the effect that they are applicable, since some conceptualization would have to precede any such judgment so as to bring into view the putative conditions for the concepts applicability.

There is no obvious reason why the conceptual structuring of our immediate perceptual apprehensions of the world should be expected to stop at the doorstep of practical or action-guiding concepts, and it doesn't. For instance, when a seasoned driver approaches an intersection with a red light, she sees the light immediately as a reason to stop, not as a visual signal that she must then *judge* to be a reason to stop. Similarly, a competent woodworker, entering her workshop with an eye to building a table, does not see undifferentiated objects whose role in her plans awaits to be determined by practical deliberation. She notices matching colors of a set of cherry planks as lending themselves to being joined in an appealing tabletop, and she sees imperfections on the edges of these boards as reason to start up the jointer and run them through.... It is our desires that continuously set the stage for practical reasoning by providing us with a pre-reflective picture of how and why it makes most sense to extend our activities. (2009: 31).

⁶⁸ Patterns of attention are one of the empirically well-established ways that higher-level thought is understood to influence perception, since it, at the very least, changes the inputs to perceptual processing (Firestone and Scholl 2016).

I quote this lengthy passage because unpacking it will demonstrate the bridge between perceptual representation and practical thought, and also illustrate how the evaluative outlook view meshes with the account of representational mind I have been developing here.

First, I must point out that Brewer's notion of perception is not the same as mine. He clearly means something like conscious perceptual experience. When he says that a neurologist "sees" dendrites, I think we can take him to be referring to visual perception, but in a way that is neutral on whether the visual system itself takes conceptual inputs. His language of "immediate apprehension of the world" suggests some sensitivity to the point that perception, as a psychological kind, is not identical to conscious perceptual experience. But even if he is not, it is clear that "immediate apprehension" refers to a representational state that plays a certain epistemic role. That is, immediate apprehensions can serve as starting points for inference and deliberation. And in order to play that role, they cannot themselves be the products of "implicit" inferences, on pain of regress. That is presumably why it is an "unstable position" to think of conceptualizations as the products of judgments.

Reflecting on the epistemic role of perception highlights the need for a practical analog to perception. On the evaluative outlook view, that practical analog is a kind of evaluative representation. Epistemology, as I understand it, concerns norms for belief and knowledge, and so only applies to individuals with propositional representational capacities—capacities for thought. In such individuals, perceptions provide warrant for perceptual beliefs. This perceptual warrant is non-inferential, and need not be reflectively accessible to the individual. Burge (2003) calls this kind of warrant *entitlement*, and he differentiates it from *justification*, a kind of warrant that is reflectively accessible to the individual. Inference, in my view, is logical (that is, truth-preserving) transformation from one conceptual or propositional representation to another.

When Brewer says that the conceptual structuring of our immediate apprehensions of the world does not depend on judgments, I understand him to mean that there is an important level of conceptual structuring that is not the result of inference, and that is not reflectively accessible to the individual. Just as we are entitled to our perceptual beliefs, we are also entitled to our evaluative beliefs formed through similar representational processes. Practical thought, just like theoretical thought, has a foundation in lower-level representational capacities.

But the lower-level capacities are “pruned and shaped” by higher-level cognitive processes, including education and rational reflection (Brewer 2009: 30). It is important for the evaluative outlook view that this be so in order for evaluative outlooks to explain dialectical activity and to provide a plausible account of rational versus irrational action. I call conceptual representations (including but not limited to beliefs) formed directly on the basis of perception *perceptual cognitions*. So evaluative representations with this kind of content are *evaluative perceptual cognitions*. When perceptual cognitions have propositional structure, they are perceptual beliefs. I have argued that many evaluative and non-evaluative perceptual cognitions are formed through similar processes, but one might reasonably question whether we are entitled to evaluative perceptual cognitions in the same way that we are entitled to non-evaluative ones. For now, I simply note that the degree to which we are entitled to our perceptual beliefs is not uniform across all perceptual cognition. The difference between perceptual cognitions and *desires* (which are evaluative perceptual cognitions with phenomenal modes of presentation) looms large. I develop these points in the next two chapters. My aim in this chapter is to establish perceptual cognition as a plausible candidate for the kind of lower-level representation that the evaluative outlook view needs to get its explanations off the ground.

Regardless of whether genuine cognitive penetration of perceptual subsystems occurs, I think various applications of conceptual attributives are sufficiently perception-like for that purpose.⁶⁹ The key points are that some conceptual representations are formed quickly, automatically, without conscious effort, and through processes that are reflectively inaccessible. They are *not* formed through inference. Moreover, some of these are like perceptions in that they apply attributives directly to particulars in a noun-phrase-like iconic format, partly in order to fix reference, in a context-bound application.

Let us consider in more detail how evaluative perceptual cognitions might be formed. Recall that on Burge's view, functional attributives are "higher-level" perceptual attributives. That is, they are always formed on the basis of "lower-level" attributes like shape, size, color, and so forth.⁷⁰ But conceptual attributions can also occur entirely on the basis of lower-level perceptual attributions. For example, one can see something as a piano entirely on the basis of perceptual attributions of shape, size, color, and so forth (Burge 2014).

If conceptual attributions can occur in context-bound singular reference through the same process as higher-level perceptual attributions, this would help explain how those conceptual attributions can be automatic, unconscious, and reflectively inaccessible, regardless of which specific representational system they occur in. And there is no reason to think that evaluative conceptual attributives could not be applied through the same processes. Thus, evaluative concepts—indeed, any concepts that can be applied to particulars in the individual's

⁶⁹ Maura Tumulty (2020) argues that for most epistemological and ethical purposes, the exact ontology of perceptual states, and the exact relation between perception and cognition are issues we can remain largely agnostic about. I tend to agree.

⁷⁰ Burge focuses on vision. But one can see, I think, the general shape of how to adapt this basic point about higher-level attributives being based on lower-level ones to other perceptual modes. Burge concentrates on vision largely because vision science is more developed than the perceptual psychology of other sensory modes; however, it may be that treating vision as the central case is limiting. See Begby (2011) for an argument to that effect.

environment—could be applied through the same kind of process that applies higher-level perceptual attributives. If Burge (2014) is correct that “one can represent something as a pushing, as an instance of agency, as a pine tree, as a piano, or as one’s favorite movie villain—*only* by visually attributing low-level attributes” then it seems likely that evaluative conceptual attributions like delicious, beautiful, or courageous could also be applied entirely through low-level attributions.

For example, consider how someone might *see* something as a piano. On this picture, low-level visual perceptual attributives for shape, size, color, and so forth would serve as a basis for that conceptual attribution. Within the visual system, those same low-level perceptual attributives serve as a basis for applying higher-level perceptual attributives, like body and food (Burge 2022: 118–120). In order visually to apply the attributive piano, I must learn what pianos are and what they look like. I could learn this at first through explicit propositional thoughts, such as “That is a piano.” But eventually the application of the attributive becomes relegated to more automatic, reflectively inaccessible processes. Perhaps these processes do not occur entirely in the visual system itself, but nonetheless they enter conscious awareness in the same way that perceptions do. And within the visual system, perceptual learning can result in the formation of an attributive piano-body. This attributive functions to pick out objects with the size, shape, color, and so forth characteristic of pianos, even though it does not function to pick out pianos as such.⁷¹ And on the basis of this kind of perceptual attribution, post-perceptual cognitive representation systems generate a conceptual representation that applies the attributive piano (which functions to pick out pianos as such). Thus, eventually I become able to have

⁷¹ It is perhaps possible that a perceptual attributive piano could exist that functions to pick out pianos as such. The key question, which is thoroughly empirical, is whether the perceptual system itself could become sensitive to the difference between pianos and non-piano but piano-like objects in the environment (Burge 2022: 120–124).

representations with noun-phrase-like content such as piano there. Such a representation applies the higher-level attributive piano on the basis of low-level visual attributives (shape, size, and so on), and it does so as part of how it picks out the object in context-bound singular reference.

Evaluative concepts get applied in the same way. To see an apple as delicious or an act as courageous, the individual must first apply lower-level perceptual attributives. Importantly, *a priori* considerations speak against the possibility of perceptual attributives for many evaluative properties—moral properties, for example—since the properties such attributives pick out are not sensible. But conceptual attributives for such properties could nonetheless be applied directly on the basis of other lower-level perceptual attributions (through learned association, for example).

Perceptual cognitions, then, are representations that apply conceptual attributives to particulars on the basis of lower-level perceptual attributions. Evaluative perceptual cognitions are those perceptual cognitions that apply evaluative conceptual attributives.⁷² Evaluative perceptual cognitions are the “immediate apprehensions” that, for Brewer, play a crucial role in rational agency. And it is essential that they can have rich conceptual content, because without it they cannot explain dialectical activity.

Perception guides agency by setting goals for perceptual-level conative states. A spider, for example, perceives a physical body in its environment as having such and such a shape, size, or pattern of movement, and on the basis of those perceptual attributions applies the functional attributive food. This attribution then leads to the perceived object becoming the goal of an

⁷² As I have intimated, I find it doubtful that perceptual cognitions occur through strict top-down penetration of conceptual attributives into perceptual subsystems. But whatever process applies conceptual attributives in this way is *not* attributable to the individual, even though the results of that process—the cognitive representational state—is attributable to the individual. This marks another feature that perceptual cognitions and perceptions have in common. Thus, I endorse Burge’s *anti-individualism*: the view that (i) representational states are partly constituted by causal relations between the individual and its environment, and (ii) that objective representation arises through processes that are not attributable to the individual. I cannot develop or argue for this view here. Burge (2010a) articulates the core arguments in favor of the view, and I find those arguments generally persuasive.

action command (perhaps for a jump) that determines the spider's movements in catching and eating that object.⁷³ Evaluative perceptual cognitions play a similar role in more sophisticated psychologies: I perceive the apple on the kitchen table as *delicious*, and so the apple becomes the goal of an action command that determines my movements as I pick up the apple and take a bite. Obviously, these examples gloss over many important differences between the actions of the spider and the human being. For example, the human being might not immediately form a perceptual-level command, but instead form an *intention* to eat the apple; the spider, since it lacks a capacity for propositional thought, cannot form intentions. Similarly, the human being might refrain from forming *any* action command at all, despite perceiving the apple as delicious, because of other propositional representations—that the apple belongs to someone else, for instance. In contrast, if the spider refrains from forming an action command to catch and eat the prey, that would be explained by more basic processes, which can be described computationally and causally; for example, if it recently fed, the spider may not be hungry. Nonetheless, these examples illustrate the general structure of how evaluative representations, both perceptual and cognitive, function to guide representationally mediated action.

IV. Desire and the Problem of Receptivity

Brewer's version of the evaluative outlook view depends on a particular notion of desire, which I have largely embraced. But there are worries about the claim that desires are "seemings of goodness" in the way that Brewer proposes and that I take up. I conclude this chapter by defending my evaluative outlook account against a general criticism by Tamar Schapiro (2014).

⁷³ I borrow this example from Burge (2010a). See especially pp. 514–518 for a discussion of empirical evidence that jumping spiders have sophisticated perceptual representational capacities that explain their hunting behavior.

My defense will, I hope, serve to demonstrate the fruitfulness of applying insights from representational psychology and philosophy of mind to ethical theory. I argue that noting some key distinctions, especially between conative states with and without conceptual content, deflates Schapiro's worries considerably.

The notion that desires are appearances of value is controversial. I have aimed to bypass theorizing specifically about desire by focusing instead on the representational mental states necessary for an adequate practical psychology, and bringing in a notion of desire only afterward, rather than starting my theoretical project with desire. Partly motivated by Brewer's arguments about the importance of explaining dialectical activity, I have proposed that there is a kind of evaluative representational state, distinct from both perception and belief, that not only enables rational explanations of action, but also *psychological* explanations. In part following Brewer, but also common usage, I have used the term "desire" to indicate this kind of state. Desires, as I use the term, are goal-setting conceptual conative states that have ineliminable phenomenal quality and that are not formed through inference. Paradigmatically, they have a perception-like format—that is, a noun-phrase-like iconic presentation in which attributives are applied partly in order to achieve reference. As goal-setting conative states, desires have evaluative representational content and motivational force.

Schapiro (2014) proposes what she calls the "paradox of passion" as a central problem that any substantive ethical theory must resolve. The paradox is this: How can we make sense of spontaneous motivations, given that such motivations must be both active, in the sense that the agent is still the source of her action, *and* passive, in the sense that such motivations are not the result of practical deliberations? This problem is clearly a formulation of what Korsgaard (2008) calls the problem of receptivity. Schapiro criticizes evaluative outlook views, including

Brewer's, on the grounds that such accounts obscure this problem and offer only hand-wavy solutions to it at best. Now, I think Schapiro overstates the force of this kind of worry for evaluative outlook accounts, but I accept the challenge nonetheless. My version of the evaluative outlook view, according to which desires are phenomenal conative states with perception-like structure and conceptual evaluative content, can answer Schapiro's worries. Or so I shall argue.

Schapiro distinguishes between two notions of desire, which she calls "placeholder" notions and "substantive" notions. Placeholder notions of desire rely on the idea that rationalizing explanations of action must appeal to something like what Brewer calls an apprehension of value. Rationalizing explanations must say something about what good the agent saw in her action, for without that element we lose the sense that we are talking about an action at all, as opposed to mere behavior. In short, a placeholder notion of desire, according to Schapiro, simply posits "desire" as whatever psychological state plays this role on the grounds that something must do so in order to distinguish between genuine actions on the one hand, and reflexes and nervous ticks on the other. Using the notion of desire in this way, Schapiro claims, does not say very much: "It seems relatively uncontroversial that an evaluative outlook has to play some role in the explanation of any action" (143). Substantive notions of desire are more specific. These notions restrict "desire" to only those appearances of value with respect to which the agent is passive. Thus, a substantive notion of desire identifies desire as an appearance of value that is not the result of active deliberation.

Schapiro's central claim is that evaluative outlook accounts of desire like Brewer's are ambiguous about which notion of desire they aim to explicate. She thinks that they help themselves to the easy point that rationalizing explanations of action need to include an agent's evaluative point of view, but then they slide into talking about desire as a specific kind of mental

state without acknowledging the difficult question of our passivity with respect to our desires. The central question that she thinks gets lost here is the question of “how we can be passive in relation to our own self-movement” (149). Her point is that even though we are in some sense responsible both for our actual actions and for our desires, we do not stand in the same normative relation to our actions as we do to our desires. The challenge, as Schapiro sees it, is to explain how we can be passive with respect to our desires without relinquishing the idea that they are indeed ours, for clearly our desires are not wholly alien, or external, forces pushing us around like billiard balls.⁷⁴ And evaluative outlook accounts as they stand she thinks fall short of this.

Part of what motivates Schapiro’s critique is the distinction between those actions that we desire wholeheartedly to do and those that we do more begrudgingly:

The fact that we sometimes act without having a desire in the substantive sense is entirely compatible with the central claim that whenever we act, we see something good, or reasonable, or attractive, or appropriate in doing what we do. When you take out the garbage even though you don’t feel like taking out the garbage, you are still guided by your conception of what makes taking out the garbage worthwhile. The question is whether this claim tells us anything about the desire that you happen to lack in this case—namely a desire, in the substantive sense, to take out the garbage. (138–139)

The more substantive explanation of the passive kind of motivating psychological state that Schapiro is asking for is exactly what I hope my account here can provide. By supposing that desires are phenomenal conative states with perception-like structure and conceptual evaluative content, we can shed light on why we are largely passive with respect to our desires. We are more passive with respect to our perception-like states in general—the processes that generate perception-like states are not attributable to us as individuals. Moreover, because our

⁷⁴ See Maura Tumulty’s (2020) book, *Alien Experience*, for an extensive exploration of this idea of being alienated from our own desires and experiences. Tumulty’s account is in many ways congenial to my view here.

conceptual states are shaped by our understanding of the concepts being applied, my account also can offer a compelling explanation of the relationship between desires and rational capacities. Experience and reflection shape our understanding of things over time, and that understanding in turn shapes our desires. Conception and understanding are constitutively rational capacities. And since desires themselves are attributable to us as individuals and (albeit indirectly) sensitive to these rational processes—processes which are also attributable to us as individuals—desires are thereby properly counted as ours.

So part of my response to Schapiro is to accept her challenge, and to argue that my notion of desire does not suffer from the shortcoming she presents. The other part of my response is to argue that she overemphasizes the significance of the gap in evaluative outlook accounts.

Schapiro addresses the point made by Brewer and other evaluative outlook theorists that desires are in many ways like perceptions. She thinks that the analogy between desires and perceptions only serves to push the question posed by the passion paradox back: “How is our passivity in relation to our perceptions compatible with the idea that those perceptions can, at least in principle, make genuinely cognitive contributions to our theoretical judgments? The theorist of perception has to confront this question, just as the theorist of desire in the substantive sense has to confront the analogous question” (145). But I think that the question seems on its face much less worrisome in the case of perception. Perceptions provide warrant for our perceptual beliefs, and in normal circumstances we are entitled to rely on that warrant, even if the grounds of those beliefs are not reflectively accessible to us (Burge 1995, 2003, 2010a). The same points apply to young children and non-human animals (at least, those non-human animals that have both capacities for perception and for thought). In the theoretical realm (assuming for the moment that a clean distinction between the theoretical and practical realm is possible) there

is a substantial set of givens with respect to which we are largely passive and receptive, but that the more active aspects of our cognition work on and refine; this picture seems for the most part unproblematic, if not simple or completely understood. Pointing out the analogy between the passive elements of theoretical and practical psychology does lower the temperature somewhat on the apparent paradox of the passions. Schapiro overemphasizes the importance of the puzzle in the theoretical realm, rather than seeing the theoretical counterpart to desire as reason to be less puzzled by the problem in the practical domain.

One might think, however, that there is a significant difference between perception and desire that lends itself to Schapiro's view. We seem to be much more likely to hold ourselves responsible for our desires than for our perceptions. Thus, even if Schapiro were to grant that the problem of receptivity in the theoretical realm is less pressing than she suggests, she could still say that it is pressing in the practical realm. Once again, I am sympathetic to her concern. We do hold ourselves and others responsible for desires, at least in some cases, in a way that we do not for perceptions. And the relation between representation and motivation, which I think is in the background of Schapiro's thinking, presents genuine philosophical puzzles. Still, leveraging some of the distinctions I have been attending to can help here.

To begin with, in my view perception is not directly analogous to desire, since desire, as I use the term, is conceptual and perception is not. But also, there is a kind of non-evaluative representation that is conceptual *and* with respect to which we are largely passive. That is, there is an analogue to desire in the theoretical domain. A perceptual cognition without phenomenal evaluative content is such a state. Such a state has conceptual content, but it is not propositionally structured. It is shaped by understanding, but it is not a product of inference.

Consider once more Brewer's example of the trained neurologist who understands neural anatomy and thus perceives images through a microscope as dendrites, ganglions, and so forth. Her perceptually enabled conceptual representations of the parts of the brain will enter her deliberations and reasonings already structured by those concepts; she will not have to deliberate about, infer, or judge in each instance that that is a dendrite and so on. It may even be that some of the relevant concepts, like dendrite, have become genuine high-level perceptual attributives within her visual system. But regardless of the exact psychological mechanism, the neurologist is clearly passive with respect to whether the brain parts appear to her as dendrites, ganglions, and the rest. But there is nonetheless an important sense in which she is responsible for those representations. The representational states themselves are attributable to her as an individual, even if the *processes* through which they are formed are attributable only to various subsystems of her psychology. As conceptual representations, they are the products of her *understanding*, which she has developed over time through deliberation, reflection, deliberate direction of attention, and intentional action.

Upon reflection, then, there is not such a great divide between the practical and theoretical domains after all. And if passivity in theoretical cognition is not a deep problem, perhaps it is not such a great concern in the practical realm either. Moreover, Brewer's evaluative outlook account of desire does offer substantive claims about the nature of desire: for Brewer, desires have evaluative conceptual content, and they are ineliminably phenomenal. Perhaps Schapiro is right not to be satisfied with what Brewer says on that score, but I think Brewer's view is on the right track.

Chapter 3

Emotions, Value, and Phenomenology

The phenomenal quality of many evaluative representations plays an important role in the evaluative outlook account of practical reason that I endorse. And for many thinkers, emotion has seemed like a natural starting point for thinking about evaluative representations and phenomenology. In this chapter, I consider how emotions might fit into my account.

Christine Tappolet (2017) argues that emotions are perceptions of value. That is, emotions are constitutively perception-like (which for Tappolet just means non-conceptual) representations of evaluative properties. In addition, she defends *Representational Neosentimentalism*: the view that evaluative concepts—paradigms include the admirable, the fearsome, and the disgusting—apply only if and because particular emotional responses are appropriate. In what follows, I argue that both of these views are problematic. The perceptual theory of emotions relies on a vague notion of perception, and once we articulate a more careful conception its plausibility diminishes. Representational Neosentimentalism fails to account for the explanatory priority of values over emotional responses; by trying to explain value in terms of emotional responses, the view fails to deliver on its promise to shed light on the nature of value. Nonetheless, both of these views contain important insights that point us toward a better understanding of emotion and value. The perceptual theory highlights the importance of evaluative representations. Neosentimentalism draws attention to the epistemic role of phenomenology in developing evaluative understanding. Taken together, these insights point

toward the possibility that phenomenal modes of presentation are necessary for representing the world of value accurately.

In this chapter, I start by discussing the perceptual theory of emotions, questioning whether emotion really could be a genuine kind of perception. I then discuss Representational Neosentimentalism, a theory about the epistemology of value that typically goes hand in hand with the perceptual theory of emotion. Here I argue that Representational Neosentimentalism by itself does not offer a good explanation of the relation between emotions and evaluative understanding, since it claims, but does not explain why, emotional experiences are necessary for evaluative understanding. This leads me to my main conclusions: 1) Emotion may be too fraught and heterogeneous a category to serve as the basis for philosophical theorizing about value and the epistemology of value; and 2) evaluative understanding probably does depend in some way on emotional experience, but what this tells us about the nature of value is not straightforward. In the last two sections, I clarify some important distinctions and explore the possibility that there are representational states with necessarily phenomenal modes of presentation, arguing that if there are such states, at least some evaluative representational states are among them.

I. A Perceptual Theory of Emotion

According to the perceptual theory of emotion, emotions just are perception-like representational mental states that pick out evaluative properties. Thus, a fear state, for example, is constituted by a representation of something as fearsome. In what follows I focus on the version of the perceptual theory presented by Tappolet (2017).⁷⁵

⁷⁵ The thesis that emotions are a species of, or can be understood by significant analogy to, perceptions has been defended by many philosophers. See, for example, Deonna (2006), Döring (2007), Elgin (2008), Prinz (2004), Wringer (2015), and Zagzebski (2003).

Tappolet argues for her version of the perceptual theory of emotion by (1) pointing out analogies between emotions and perceptions and (2) highlighting limitations of other competing theories of emotion (such as desire theories and judgment theories). Here I focus on Tappolet's positive argument for the perceptual theory.

Emotion, Tappolet claims, is not identical to sense perception, but is like it in important ways, and also unlike it in important ways. Emotions are analogous to perceptions, she says, in that they are (a) paradigmatically conscious states characterized by phenomenal properties, (b) not directly subject to the will, (c) typically caused by events and states of affairs in the world, (d) assessed according to their correctness conditions, (e) informationally encapsulated, (f) not incorporated into inferential networks, and (g) composed of analog (as opposed to digital) content (2017: 19–24). Emotions are unlike perceptions, though, in that they are (a) not dependent upon specific sensory organs, (b) more phenomenologically complex, (c) characterized by positive or negative valence, (d) not as transparently about mind-independent objects, (e) affected by socio-cultural factors and changes in the individual's psychology over time, and (f) subject to rational constraints (24–28).

Despite these impressive differences, Tappolet argues, the perceptual theory of emotions is still a viable theory as long as we adopt a sufficiently capacious understanding of what perception is. She suggests analyzing perception as a general “openness to the world,” understood as direct awareness of the world and its properties (29–30). Emotions and perceptions play the same kind of epistemic role, according to the perceptual theory, in that they provide foundational justification for a related class of judgments. Just as our basic judgments about the physical world are based on sensory perceptions, our basic judgments about the evaluative world are based on our emotions.

One problem with this view is that it relies on an empirically implausible account of perception.⁷⁶ On the most empirically plausible account of perception, perception is not direct awareness of the world; rather, it involves the application of perceptual attributives to particulars. Perception, like all forms of representation, is a kind of mental perspective on objective reality. On any straightforward understanding of what direct awareness means, perception is not direct awareness. Rather, perception is awareness mediated by representational systems that are *causally* connected to their subject matter. This causal connection partly explains how perception can play the epistemic role that it does. If a similar causal relationship does not exist between emotions and their subject matter, then it is not clear how emotions in general can play the same kind of epistemic role as perceptions. And pretheoretically, emotion is a diverse category. It may be so heterogenous that it resists a cohesive theoretical treatment. Indeed, Tappolet claims that the category of emotion may not be a natural kind, but that “emotions form a unified category” nonetheless (7). But this suggestion is unhelpful as a strategy to defend the perceptual theory, since perception *is* a natural kind. If emotions are not a natural kind, then the claim that emotions are perceptions of value looks much less plausible, regardless of how useful the category of emotion turns out to be. Thus, to mount a rigorous defense of the perceptual theory of emotions, one would have to show that emotions form a natural kind on the same order as perceptions, or that emotions form a natural sub-category of perceptions.

Moreover, perception does not seem to depend, as a constitutive matter, on phenomenal consciousness. But emotions are constitutively conscious states.⁷⁷ There is good reason to think that relatively simple organisms have representational states, perhaps including states with

⁷⁶ See Burge (2010a, 2022), as well as the previous chapter of this dissertation.

⁷⁷ This characterization of emotions is not uncontroversial. For example, Burge (2022) does not seem to agree. See *Section IV: Presentation and Phenomenal Consciousness* below.

representational content similar to that of some emotions, but it is unknown whether these organisms are phenomenally conscious (Burge 2010a: 375). An organism without phenomenal consciousness cannot have mental states with phenomenal qualities, and so cannot have emotions. Insects, for example, clearly have capacities for perception. But it is not clear whether insects are conscious. Whether or not any insects have capacities for phenomenal consciousness, the question of insect consciousness comes apart from the question of insect perception. This epistemic point suggests a significant disanalogy between perception and emotion.

Even in humans, there are representational states that lack conscious presentation (Burge 2010a: 374–375). Plausibly, unconscious states are at least sometimes capable of applying functional attributives like danger. A perceptual theory of emotion, which types emotions by their representational content, is ill-suited to account for the distinction between non-phenomenal perceptual states that apply functional attributives (like unconscious perceptual representations as of danger) and phenomenally conscious states that apply those same attributives (like fear).

So having in view a substantive and empirically plausible account of perception makes a perceptual theory of emotions complicated to defend. Tappolet seems to be somewhat sensitive to this problem, since she suggests that anyone opposed to her “liberal account of perception” might still accept that emotions are at least “quasi-perceptions” (2017: 30–31). This amounts to a mere terminological dispute, she claims, since the important point is that emotions are “states that have representational content, so that they can be said to be correct depending on whether their objects have the corresponding evaluative properties or not, but that are nonetheless different from evaluative judgments” (31). The point of the perceptual theory of emotion, then, is to emphasize that emotions are representational, but that their content is non-conceptual.

If we grant this point to Tappolet, we are left with a claim that there is an important kind of mental state, call it emotion, that has non-conceptual evaluative representational content. Moreover, Tappolet makes the following important qualification: “Although emotions can, and often do, involve conceptually articulated contents... the evaluative appraisal that is part of the content of emotions is non-conceptual” (18).⁷⁸ So the key task for any theory of emotion aiming to go in this direction is to articulate the nature of the non-conceptual evaluative content that emotions supposedly have. And the nature of this content will also ideally help explain the epistemic role of emotions. After all, this epistemic point is supposed to be a primary virtue of the perceptual theory: emotions are to evaluative judgments as sense perceptions are to judgments about our physical environment. I think one of the most promising ways to cash out this idea is to suppose that the non-conceptual evaluative content of an emotional representation resides in its phenomenal quality. This suggestion has the advantage of tying together the two key constitutive elements of the kind of mental state we are trying to understand: their representational content and their phenomenology. The idea is that the phenomenal quality of the state *is* the mode of presentation of the evaluative element of its content. Before I discuss this suggestion, however, it will be helpful to consider the theory of value that makes up the other major strand of Tappolet’s view: Representational Neosentimentalism.

⁷⁸ Conceptual content, Tappolet points out, is posited in order to account for inferential networks. Mental content that is incorporated into an inferential network, the thinking goes, must have repeatable components that occur in all the representational states that are inferentially connected. It is this repeatable content that accounts for the inferential connections.

II. An Emotional Theory of Value

Representational Neosentimentalism claims that all evaluative concepts can be explained in terms of emotional concepts. Tappolet (2017) presents the view in terms of the following biconditional:

(R-NS) x is V if and only if x is such that feeling E is correct in response to x , if one were to contemplate x .

Here V is an “affective” value, such as admirable, and E is the corresponding emotional attitude, such as admiration.⁷⁹ So, putting the idea in (slightly) plainer language, something is admirable (or shameful, fearsome, or whatever) if and only if feeling admiration (or shame, fear, and so on) is a correct response upon contemplation of that thing.

Now, this biconditional is doubtless true, at least if we take it at face value. But it is not clear whether it is important or interesting. What it needs, I think, is some substantive account of how an emotional state, like admiration, sheds light on a genuine evaluative property. After all, offering necessary and sufficient conditions is not the same as an explanation. The most obvious version of such an account is reductive: if evaluative properties just are (perhaps disjunctive) non-evaluative properties that happen to (reliably, dispositionally, or typically) generate emotional states in human beings, then R-NS could be understood as a statement of that view, and thus as saying something interesting about the nature of value. But then the obvious truth of the biconditional would evaporate, and it would take on a heavy burden of argument. No wonder,

⁷⁹ Tappolet does not offer a precise definition of *affect*. She suggests that affect is a similar but more general notion than emotion (2017: 178–180). And she uses the term “affective concept” to refer to evaluative concepts that are tied to specific responses such as admiration, fear, disgust, and so forth. Since she takes these responses to be clear examples of emotions, I interpret her use of the term “affect” as a label for the phenomenal experiences characteristic of emotions. This notion is more general than emotions, since emotions are representational states, and affective states need not be representational.

then, that Tappolet rejects reductive versions of Representational Neosentimentalism (2017: 83–90).

On Tappolet’s view, the biconditional makes an *epistemic* claim. The idea is that emotional experiences are indispensable for understanding evaluative concepts like admirable, shameful, and fearsome. Representational Neosentimentalism “aims at explaining the *possession conditions* of evaluative concepts, something which is claimed to involve experiencing emotional responses and grounding one’s evaluative judgments on these responses” (100). In Tappolet’s view, the epistemic claim made by Representational Neosentimentalism is twofold: On the one hand, emotional experiences are *developmentally* indispensable to our evaluative understanding; on the other, emotional experiences are essential to our epistemic practices of making particular evaluative judgments. And these epistemic claims mean, she thinks, that emotional concepts should be the basis of our philosophical explanations of the *ontological nature* of value (2017: 116–121).⁸⁰ For this inference to work, it must be the case that affective emotional concepts are the most basic or fundamental kind of evaluative concepts. Tappolet argues that they are.

Tappolet thinks that evaluative concepts in general acquire their evaluative content through some kind of connection with emotional experiences. For example, she points out that courage, a so-called “thick” evaluative concept, is always admirable (111). In addition, emotions often have positive and negative valence, which corresponds with “thin” evaluative concepts like good and bad (113–114). If this view is correct, then affective emotional experiences are epistemically indispensable for understanding value *in general*. On this view, evaluative concepts must ultimately be explained in terms of emotional responses; moreover, emotions must be explained in terms of evaluative concepts. For instance, fear is constituted in part by its formal

⁸⁰ Tappolet emphasizes, however, that representational neosentimentalism is compatible with many different ontological theories of value (116).

object, the fearsome. Thus, according to Tappolet's view, fear cannot be explained without reference to the fearsome, since that evaluative property makes up the veridicality conditions constitutive of fear. But fearsomeness also cannot be explained without reference to the emotional experience of fear.

The problem for Tappolet's view is that it ends up with an unnecessary philosophical quietism about value. Tappolet claims that Representational Neosentimentalism is not a theory about the application conditions of evaluative concepts because we cannot know those application conditions *a priori* (99–100). Representational Neosentimentalism is just a theoretical “framework” reminding us that if we want to understand some evaluative property, then “we have to start with our affective responses” (102–103). And the lesson seems to be that we just have to be comfortable with some fairly tight circularity in our philosophical theorizing. For example, we might ask what the emotion of admiration is. The neosentimentalist answers that admiration is an emotional response that constitutively picks out (functions to represent) the evaluative property admirability. And what is admirability? Admirability is the property to which admiration is a fitting response. We should not, Tappolet insists, ask for more from our philosophical theory. The idea, it seems, is that if we want to understand admirability and admiration better, we will need to reflect on the various specific things that prompt us to feel admiration, and perhaps also on the qualitative experience of admiration itself.⁸¹ That is a point well taken. But I think theory can do more, both by providing direction for our further reflections and by explaining *why* reflection on specific emotional experiences is epistemically necessary. The concept of admirability, for example, has more content than “that for which admiration is a

⁸¹ It is worth pointing out, however, that Tappolet does not emphasize the specific phenomenal qualities of emotions, and so does not say anything about the potential fruits of reflection on those qualities. In what follows, I attempt to fill some of this gap.

fitting response.” There is a lot we can say about the nature of admirability that goes beyond the appropriateness of admiration, but falls short of a complete list of application conditions. A mature evaluative understanding, even if it causally and developmentally depends on emotional experiences, ultimately comes free from those experiences on a conceptual level. Such a mature understanding of value, therefore, can be brought to bear in explaining emotions in a way that an understanding of emotions cannot be used to explain values. In short, I think a correct theory of value will not ground the nature of value as such in emotional responses, and thus it will be in a position to explain the epistemic role of emotions in more depth than Representational Neosentimentalism attempts to do.

Before I say more, I want to get a better grip on the essential problem with Tappolet’s account. The basic problem, I think, is that Representational Neosentimentalism, as she presents it, gets the order of explanation between emotions and values wrong.

For example, in her effort to analyze thin evaluative concepts like good and bad in terms of emotion concepts, Tappolet suggests that thin evaluations like good and bad can be understood as referring to those things that warrant positive and negative emotions respectively. Thus, she presents yet another biconditional:

(R-NS-good) x is good/bad in a way iff x is such that feeling a positive/negative emotion is correct in response to x , if one were to contemplate x .

This is meant to shed light on both thin evaluative concepts and emotion concepts. But I think it does neither. When she considers the obvious question of what makes an emotion positive or negative, the answer she settles on is that positive emotions are tied to positive evaluative concepts, and negative emotions are tied to negative ones (114). And she shrugs off the glaring circularity problem: “if the goal were reduction, this would make far too tight a circle.

But in so far as what is aimed at is elucidation, this kind of circularity is unproblematic” (114). This answer to the threat of circularity is much too quick.

First, the biconditional R-NS-good cannot be understood as an epistemic claim in the way that the previous biconditional R-NS might plausibly be. There are straightforward concepts of goodness that can quite easily be understood without any reference to emotions. Take for example goodness relative to function. Getting enough water is good for a tree. This notion of goodness can be explained in terms of various biological and teleological facts about the tree. No reference to any emotions is needed.

More to the point, the question of explanatory primacy is left untouched. If “elucidation” just means pointing out that there is *some* dependency relation between two things, then it is of little value. A satisfactory explanation ought at least to indicate the direction of dependence. In this case, I think that positive and negative emotions are probably best explained in terms of notions of good and bad, rather than the other way around.

Consider also Tappolet’s analysis of thick evaluative concepts, like courage, in terms of emotion concepts. This analysis, while more plausible on its face than her sentimentalist treatment of thin concepts, founders on the same rock. Once again, the question of direction of explanation is unresolved. Is an act courageous because it is admirable, or is it admirable because it is courageous? Tappolet suggests that neither the emotion concept nor the thick evaluative concept is explanatorily primary (110–113). She thinks that a thick evaluative concept like courageous can be separated into a descriptive concept and an evaluative concept, along with a grounding relation. Thus, a courageous act is performed in spite of danger (description) and is admirable (evaluation) *in virtue of* precisely how it is performed in spite of danger (112–113). There are some relatively minor difficulties with this suggestion. For one, it may be that some

acts performed in spite of danger are admirable but not courageous. One might admire the wonton bravery of a villain, for example, but it is not clear that such admiration is unwarranted or incorrect, and it is also unclear that such bravery amounts to courage. For another, it seems likely that the unspecified ways that an act might be performed in spite of danger so as to ground the act's admirability are capturable with other, non-emotional, evaluative notions. For example, suppose that a courageous act is one performed in spite of danger and in pursuit of a worthy goal (or at least, a goal that reasonably seems worthy from the perspective of the agent). Such an act would surely be admirable, but that admirability would not explain the act's courageousness. Rather, the act's courageousness would be explained by the worthiness of the goal and the kind of danger faced in its pursuit.

And this brings us to the core problem. Pretheoretically, that an act is courageous explains why it is admirable, while the reverse does not hold. Thus, even if the developmental part of Tappolet's story turns out to be correct—that is, even if it is true that we come to understand evaluative concepts through having emotional experiences—the ontological part of her theory, according to which there is no explanatory priority between emotions and values, fails. In the case of courage, suppose for the sake of argument that we always come to understand the difference between courage and cowardice or rashness by experiencing admiration. This would mean that my understanding of the concept of courage depends upon my having had experiences of admiration for courageous acts. This epistemic claim says little on its own about the nature of courage. If anything, it tells us something about the nature of admiration: admiration functions to pick out agents and acts with valuable properties like courageousness.

Consider how sensory perception works in a similar way. Perceptions of physical bodies, for example, are epistemically indispensable to our conceptual understanding of physical bodies.

We acquire our concepts for physical bodies through having perceptions of physical bodies. But the ontology of physical bodies has nothing to do with sense perceptions and on a conceptual level is explained entirely independently of them. Rather, we understand (in a theoretical way) the nature of sense perceptions of physical bodies in terms of our independent understanding of the physical bodies themselves: perceptions as of physical bodies function constitutively to pick out those things in our environment that are physical bodies. Here we find an important principle about the relation between epistemology and ontology: If understanding a subject matter is epistemically dependent on a kind of mental state, that does not entail that the subject matter is ontologically dependent on the nature of that kind of mental state. It only entails that the subject matter must have a nature such that it be representable by mental states of that kind.

Tappolet recognizes this principle, claiming that the only ontological constraint on the nature of evaluative properties posited by R-NS is that evaluative properties must have a nature such that they can be represented by perception-like emotional states (2017: 119). But because her notion of perception is so abstract and capacious, that constraint amounts to almost nothing. Once we have a more precise understanding of perception, and of representation in general, we can apply this constraint more usefully by considering the various modes of presentation involved in representations of values and evaluative properties. Insofar as sensory perceptions tell us anything about the nature of the things they pick out, we must look at the specific content of the perceptions. Whatever the natures of the entities represented in sense perceptions, those natures must be such that sense perceptions are capable of representing them *in the modes of presentation that the sense perceptions have*. If the analogy between sense perception and emotion is to hold, then, we should expect that insofar as emotions tell us anything about evaluative properties, we will have to look at the specific modes of presentation of those

emotions. But Tappolet says nothing about the modes of presentation of emotions. In what follows, I try to fill this gap.

III. The Nature and Epistemic Role of Emotions

Let us consider more closely the epistemological picture that Tappolet presents. This is where her view is most clearly on the right track. The first key point is developmental. We acquire a mature evaluative understanding through a process of experiencing various emotions. On this view, our evaluative understanding starts with emotional evaluative concepts like admirable, fearsome, shameful, disgusting, and the rest. In this way, there is always a causal and developmental link between emotions and evaluative understanding.

The second component of her sentimentalist account of evaluative understanding is a claim about our actual epistemic practice of making particular evaluative judgments. According to Tappolet, emotions are to evaluative judgments as sense perceptions are to judgments about our physical environment. Thus, my fear of the dog provides *pro tanto* justification for my belief that the dog is fearsome. This second component of the sentimentalist view is connected to the first. Experiencing emotions and grounding evaluative judgments on those emotions is part of the process by which we develop our evaluative concepts and, consequently, our evaluative knowledge and understanding.

A third component of the sentimentalist view, which I have not yet discussed, is the role of “comparative discounting practices” by which we make sense of inter- and intrapersonal discrepancies in our emotional responses to things. These discounting practices include all the ways that we correct, check, and refine our emotional responses by comparing them to other

similar responses, both our own responses and those of others. Pointing to these practices is part of how Tappolet attempts to deal with the apparent circularity of her account (2017: 101–102). The threat of circularity arises because her account does not provide substantive content to evaluative emotional concepts (like fearsome) beyond that the corresponding emotion (fear) is an appropriate response to things that possess the property picked out by the concept. Discounting practices, she suggests, can help add more specific content to those evaluative emotional concepts by specifying more precisely when an emotion is an appropriate response, even though the specific conditions fit to survive our discounting practices cannot be specified *a priori* (101–103, 170–172). Discounting practices also play a role in the epistemological function of emotions, according to Tappolet, since discounting practices essentially aim at discovering defeaters that undermine the justificatory power of emotions (170–172). Thus, on her view, emotions provide justification for evaluative beliefs as long as no defeaters are present, just as sense perceptions provide justification for beliefs about the physical environment as long as no defeaters are present. In both cases, our comparative discounting practices aim at identifying defeater conditions and thus shoring up shortcomings in our capacities for representing reality as it is.

Now, this general epistemological picture has a lot to recommend it. I think Tappolet is correct that emotional responses play a role in the development of evaluative understanding. Moreover, we certainly do form evaluative beliefs on the basis of our emotions, and those beliefs are at least sometimes warranted. Finally, I share her doubts that a perfectly general theory of value can be formulated without any reference to our specific emotional responses and the circumstances that give rise to them. But there are some problems with the account as it stands.

First, it is misleading to say that emotions and sense perceptions *justify* related classes of belief. There are different kinds of warrant. A belief warranted by reflectively accessible inference from other beliefs is *justified*. Sense perceptions warrant belief through *entitlement* (Burge 2003). This is more than a mere terminological point. The distinction between justification and entitlement serves to emphasize the different psychological processes that result in belief, processes that constitutively take different kinds of mental state as inputs. So the question that arises is whether we are *entitled* to our evaluative beliefs based on emotions in the same way we are entitled to our perceptual beliefs based on sense perceptions. Once we employ this distinction, it is a little less clear that emotions and sense perceptions have the same relation to beliefs and judgments.

The other problem with the sentimentalist account of evaluative epistemology, which is more pressing, is that emotions are not the only kind of mental state through which we grasp the world of value. Another relevant category is qualitative phenomenal states in general, a category that includes, but is much broader than, emotion. This category includes non-representational states like pleasure and pain, emotions like fear and anger, and other harder-to-label attitudes like respect and appreciation. Once we have such a wide array of psychological states in view, the neosentimentalist story about evaluative epistemology looks much less tidy.

For one thing, pleasure and pain seem to play a similar role in the development of our evaluative understanding as emotions do. One way we learn about the good and bad, at least in the early stages of development, is by experiencing pleasure and pain. But pleasure and pain are not constitutively representational, so any evaluative concepts formed on their basis will not have the same kind of possession conditions as concepts formed from non-conceptual representational states.

Moreover, focusing on phenomenal states in general brings attention to a central feature of emotions: their felt phenomenal quality. As I have argued, emotion types are primarily defined by their felt phenomenology. Fear, for example, is typed by how it feels, not just its representational contents.⁸² Yet fear is representational. And I have already discussed primitive representational states, at the pre-conceptual level of perception, that seem to have the kind of representational content associated with fear: representations that apply the functional attributive danger. These considerations suggest some important things.

First, the distinction between primitive emotions—those based on very primitive functional representational states, like fear—and more sophisticated emotions—those, like admiration, whose representational content is not clearly traceable to primitive functional representations—is more important than neosentimentalists admit. So sentimentalists have much more mapping of the territory to do than initially seems to be the case. It may be that not all emotions have the same relation between their phenomenal and representational elements. For example, it might be the case that the qualitative phenomenology of an emotion like fear is a contingent effect in creatures like us that arises alongside representations with the kind of content that fear has. That is, creatures like us might always experience the feelings we associate with fear when our primitive representational systems apply the functional attributive danger in such a way that our lower-level conative systems become primed for avoidance behavior. But those phenomenal experiences might not be constitutively connected to that representational content. The experience may just be the phenomenal side effect of the priming of the conative system, for instance, but in a creature without any phenomenal consciousness, the same representational content and conative priming would occur without any such feeling. Indeed, in the case of fear

⁸² As I discuss later, it may be that, contrary to what I say here, not all emotions are constitutively felt or phenomenal.

this analysis looks especially viable, since there are creatures that clearly have perceptual psychology, presumably including representations that apply the functional attributive danger and the attendant conative representations for avoidance, but that may well lack phenomenal consciousness.⁸³

On the other hand, the phenomenal aspect of some emotions could have a necessary connection to the representational content. It could be, for instance, that the phenomenal quality of an emotion is partly constitutive of the representational content, perhaps as the *mode of presentation* of that content. This could be similar to the way that the mode of presentation of visual perceptions of color is partly constituted by specific phenomenal qualities—the phenomenal quality of what it is like to see red, for example.⁸⁴ So an emotion like admiration, according to this suggestion, represents its object as admirable, and that representational content is subjectively presented through a particular *feeling* toward that object.⁸⁵ If this is right, an emotion like admiration cannot represent its object as admirable without also involving its characteristic phenomenal quality because that quality is *the way the representational state presents its content*.

It may be that some emotions represent their objects with constitutively phenomenal modes of presentation, while others do not. Perhaps admiration, for example, involves a constitutively phenomenal mode of presentation, while fear only contingently presents with a characteristic phenomenal quality. Thus, the difference between primitive emotions like fear and

⁸³ See Burge (2010a: 374–376). The point is not that we know there are creatures who lack consciousness but nonetheless apply perceptual functional attributives. Rather, the point is that we know creatures like insects and spiders have sophisticated perceptual psychology, which probably includes attributives like danger, but we do not thereby also know that these creatures have conscious phenomenology.

⁸⁴ See *Section V: Necessarily Phenomenal Modes of Presentation* below.

⁸⁵ I use the word “feeling” here loosely to indicate phenomenal quality, not in the technical sense that I develop below.

more sophisticated emotions like admiration might run deep. Since primitive emotions have their basis in psychological capacities that may have evolved before phenomenal consciousness evolved, perhaps the phenomenal elements of those primitive emotions have a different relation to their representational content than the phenomenal elements of representational capacities that evolved after (or in tandem with) phenomenal consciousness. This consideration suggests that we perhaps ought to reject the claim that all emotions are constitutively conscious states.

I also question the idea that all emotions represent value non-conceptually. If we understand emotions as affective states with representational content, then I see no reason to think that such states must always and only have non-conceptual evaluative content. Emotion may be a more primitive capacity than conceptual thought, but that does not entail that the evaluative element of all emotions must necessarily be non-conceptual. The key question is what role the evaluative element plays in the representational and conative psychology of the individual. I think emotion theorists are right to point out that primitive emotions seem to have non-conceptual evaluative content, and that the evaluative element can help explain the motivational efficacy of emotion. But it would be a mistake to assume that the motivating evaluative element of all representations must be non-conceptual.

Once again, the heterogeneity of emotions proves difficult to account for. Our pretheoretical notions of some emotions might turn out to be gerrymandered categories, held together by mostly superficial similarities. It may be that our pretheoretical understanding of fear, for instance, includes both primitive representational states that apply perception-level attributives for danger and more sophisticated capacities that apply evaluative attributives to conceptually articulated representations of entities, states of affairs, and events. The primitive attributive danger functions to be applied to threats in the immediate environment and to be a

goal within low-level conative states (like action schemas and motor commands) for avoidance. A more sophisticated capacity—call it *conceptual fear*—functions to apply a negative evaluative attributive to states of affairs, events, and entities that may not be immediately present in the individual's environment, and as such this capacity depends on a capacity for conceptual thought. Conceptual fear, understood in this way, functions to be a goal for a similar, but more sophisticated, kind of conative representation: intention. For example, if I am fearful about some potential future state of affairs, I will likely form intentions to act to prevent that state of affairs from coming to be. The similarities between primitive fear and conceptual fear are fairly obvious, but they may exist only at a high level of abstraction. Both kinds of state involve negative appraisals that function to be goals in actional representations for avoidance, broadly construed. But the kinds of avoidance involved in each case are very different. One is physical avoidance of an immediate environmental threat; the other is avoidance in the sense of preventing some not-yet-actual event or state of affairs. The felt phenomenal quality may be similar between primitive fear and conceptual fear (in humans at least), prompting us to think of them as the same kind of state; but the phenomenal experiences may not actually be qualitatively identical, and the relations in the two kinds of state between their phenomenal qualities and their representational contents could be very different.

These reflections leave us with a couple of noteworthy conclusions. First, our pretheoretical concept of emotion may be too multifarious to serve as a good starting point for philosophical theorizing about practical reason. I think we have a better grip on the notions of representation and felt phenomenology, and we can go a long way by reflecting on how those two psychological kinds interact. The idea of representational states with constitutively phenomenal modes of presentation is a fruitful notion for theorizing. It may be reasonable

enough to call such states emotions, but we risk importing pretheoretical assumptions that, in this case, could lead us astray. And it may be that upon reflection emotions are not constitutively phenomenal states.⁸⁶

Second, the epistemic role of emotions is probably not well-captured by a straight-forward analogy with perception. The main problem with such an analogy is that it obscures the more complex relationship between emotions and propositional thought. This complexity is partly due to the heterogeneity of emotions. Some emotions seem to be very primitive, on the same phylogenetic level as perception, which makes the analogy seem very attractive. But other emotions seem to depend on a capacity for conceptual thought, even if it is true that the evaluative element of their representational content is non-conceptual (which I doubt). Thus, to say that emotions in general provide warrant for evaluative judgments in the same way that perceptions provide warrant for judgments about one's immediate environment is likely too simple.

Nonetheless, I propose that we take seriously the idea that phenomenal states have an important, perhaps ineliminable, epistemic role to play in our evaluative understanding. Moreover, the role of phenomenology in the epistemology of value may tell us a lot about the ontological nature of value, though I doubt this insight will come in the form of any kind of reduction. In the next two sections, I take up the working hypothesis that at least some evaluative representations have phenomenal modes of presentation, and that a phenomenal mode of presentation is sometimes necessary for a state to represent the value or evaluative property that it does.

⁸⁶ See Burge (2022: 507–508 and 546, fn. 771). I briefly discuss the possibility that emotions are not constitutively phenomenal in the next section.

IV. Presentation and Phenomenal Consciousness

In order to proceed, I need to clarify some distinctions. First, we must distinguish between *presentation* and *representation*. I have already discussed representation at some length. Representations are mental perspectives on an objective reality. They involve reference and attribution. Attribution is the application of attributives, which function to be applied to particulars in order to sort (or characterize) those particulars according to their properties, kinds, and relations. A *presentation*, on the other hand, is an *immediate instantiation of subjective mental content*. This requires some unpacking.

First, a presentation is an *immediate instantiation* in the sense that it is not, in itself, mediated by any perspective or way of “seeing.” A presentation can, however, partly *constitute* a perspective. Indeed, representations, which are essentially perspectival, are always composed of one or more presentations. The same representational content can be *presented* in different ways. Thus, representational mental states are type and token identified partly by their objective representational content—the structure of the state, the attributives it applies, and the objects it refers to—and partly by their *mode of presentation*—the way the (objective) content is (subjectively) instantiated in the individual’s mind. For example, a visual perception of a small sphere and a tactile perception of the same object both apply the same perceptual attributives for shape and size, and they both refer to the same particular. But they have different modes of presentation. The subjective instantiation of that content is different.

Presentations are *subjective* in the sense that they just are mental contents of an individual subject. The notion of presentation (and by extension, the notion of representation as well) presupposes the notion of an individual. Whereas representation relates the individual to an independent reality through the mental states of the individual, presentation simply relates an

individual to its own mental states.⁸⁷ Thus, the nature of a representation partly depends on the nature of its external subject matter; the nature of a presentation depends only on the internal nature of the mental content being presented.

All mental states are presentations. When a presentation constitutes a perspective on an objective reality, it is also representational. Representations are by their nature perspectival, since they are also partly constituted by external veridicality conditions, and thus are partly constituted by relations between an individual and the world. Presentations as such are not constitutively perspectival. That is, constituting a perspective on reality is *not* part of a presentation's *nature as a presentation*. A representation is composed of a subjective instantiation of mental content (a presentation), an objective reality (a subject matter), and a functional relationship between the presentation and the subject matter (veridicality). Veridicality takes the form of either an *iconic* or a *conventional* relation (or a combination of both kinds of relation) between the presentation and the subject matter. Iconic relations are natural correspondences between a representation's mode of presentation and its subject matter; conventional relations are arbitrary relations based on agreement or past practice (Burge 2022, 296–304). I discuss these relations in the next section.

It will be helpful to consider some examples. Pains are presentations. A pain state presents itself to the subject. It can carry information—about bodily damage, perhaps—but it need not do so in order to be the kind of state that it is. Whether or not a mental state is a pain

⁸⁷ The reality represented by an individual is independent in the following sense: the nature of an entity represented does not depend on the fact that it is represented or the way it is represented. Thus, even when a subject represents its own mental states, the nature of the mental states represented depends only indirectly on the mental states of the subject. For example, I have a belief that dogs are mammals. I also have a belief that I believe that dogs are mammals. The second belief has one of my representational mental states as its subject matter, and thus the nature of its subject matter is mind-dependent in one sense. But the nature of that subject matter is still independent of the mental state that represents it. Even if I did not have the meta-belief that I believe dogs are mammals, I could still believe that dogs are mammals, and *that* mental state would have the nature it has regardless of whether it is represented or not. Self-referential representations are another matter, so I set them aside for now.

depends on the phenomenal quality of the state. (If unfelt pains are possible, their nature is still determined by the phenomenal quality they would have if they were felt.) Thus, when I feel a pain, that mental state is instantiated directly, *through* being felt. How it feels is a constitutive element of its nature as a pain.⁸⁸

A non-perceptual conscious sensation, such as the smell of roses, is a presentation. It need not have any representational function at all, and yet it is still subjectively instantiated in the individual's mind. Like a pain state, a non-perceptual conscious sensation might carry information. It might even function to carry information. But it need not do so in order to be a presentation of the type that it is. Unlike pain states, such sensations are not felt, though they do have a distinct phenomenal quality.

Seeing red is a presentation. Typically, seeing red is the mode of presentation in visual representations as of red objects. (There may be exceptions, as when one sees splotches of redness across the visual field after rubbing one's eyes. These splotches are instances of seeing red, but they may not be part of any representation as of a red object.) Like pain, seeing red is constitutively phenomenal; unlike pain, seeing red is not constitutively felt.

Perceiving something as a physical body is a presentation. It is also a representation, since it applies the attributive body to a particular. But the attributive body is itself directly instantiated in the individual's mind.⁸⁹ It corresponds to reality in that it functions to refer to entities with a certain nature—that is, physical bodies. Thus, the application of the attributive comes along with veridicality conditions. Moreover, the attributive is always applied in a specific granular way. The object is represented as a body by being presented at such and such an angle,

⁸⁸ It may be the only constitutive element of its nature. If it is, then we have the makings of an argument against physicalism, but of course, the physicalist would not concede that point, so it is dialectically moot.

⁸⁹ I have some doubts about this point related to uncertainty about the relation between presentation and consciousness. But such a digression must wait for another occasion.

from such and such a distance, and so forth. These kinds of relations between the individual and the object of the representational state are also part of the content of the representational state. And these mental contents, which relate the subject representationally to the object, are themselves directly instantiated in the mind of the subject. They are presented. They might be presented phenomenally, as when a human being visually perceives (and focuses her attention on) a physical body in the typical way. But the phenomenal quality of the presentation, if it has one, is not part of the state's nature as a perception. The phenomenal quality is part of a state's nature as a *conscious* perception, however, and that phenomenal quality does partly token-identify the state. A conscious perception of a physical body is not felt, even if it does have a distinct phenomenal quality.

The belief that animal is a dog is a presentation. It is also a representation, since it, like the previous example, functions to refer and apply attributives to a particular object (the dog), and thus has veridicality conditions (in this case truth conditions, since the representation has propositional structure and conceptual content). Also, like the previous example, the state is not constituted as the kind of thing that it is (a belief) by its phenomenal quality, if it has one. The phenomenal quality, if it has one, is part of the state's nature as a *conscious* belief, if it is conscious, and that phenomenal quality does partly token-identify the state. But the state need not have any phenomenal quality at all in order to be a belief, nor need it be conscious. Whatever phenomenal quality it may have is not felt. The mode of presentation of propositional representations is a difficult topic, which I shall, for the time being, postpone. Nonetheless, beliefs and other propositional mental states are directly instantiated in the mind of the subject in some particular way (mode), and so they are presentations. Once again, all representations involve presentations.

Conscious states occur when mental content is presented *to* the individual.⁹⁰ Conscious states can be divided into *phenomenal* and *non-phenomenal* states. Phenomenal states have essentially qualitative elements. That is, they are the kind of state that they are partly in virtue of their phenomenal quality. Pain states are phenomenal. A conscious perception as of a red sphere is phenomenal. Non-phenomenal conscious states are those that either have no qualitative elements, or have only inessential qualitative elements. Such a state is not the kind of state that it is in virtue of its phenomenal quality, if it has one. The conscious thought dogs are mammals is a non-phenomenal conscious state.⁹¹

Phenomenal states can be further subdivided into *felt* and *unfelt* states. Felt states are phenomenal states whose qualitative elements are experienced as bodily sensations. Some felt states are what I call *phenomenal affective states*. I follow Burge (2022) in using the term *affect* to refer to bodily agitation or “reaction in response to practically relevant... information.”⁹² (This is in contrast to Tappolet’s use of the term *affect*. As far as I can tell, Tappolet has no specification of the term, but she seems to use the term to refer only to conscious states with phenomenal quality.) Thus, when an affective state is phenomenally conscious, it is felt. Pain

⁹⁰ I tend to call conscious individuals “subjects,” though I have not developed any strict definition of that term.

⁹¹ I am tempted to say that in order to be occurrently conscious, the state must have some sort of phenomenal quality, even if that phenomenal quality is not part of the state’s nature as the kind of state that it is. But it could be that there is a kind of subjective awareness with no phenomenally qualitative aspects. If there are such non-phenomenal conscious states, abstract beliefs such as 2 + 2 = 4 and dogs are mammals seem like the most obvious candidates. The reason I am tempted to deny that such states could lack all phenomenal quality and yet nonetheless be occurrently conscious is that I have trouble getting a grip on what it means for a state to be occurrently conscious aside from its having a presentation with some phenomenal quality. In other words, having a phenomenal quality is part of how I understand the notion of subjective awareness. But I am far from certain this understanding is correct, and so these speculations are quite tentative.

⁹² See Burge (2022: 546, fn. 771). Burge uses the term “visceral” rather than my term “bodily,” perhaps to emphasize that the reaction in question is centrally located and diffuse throughout the organism’s body. Importantly, this notion of affect is more primitive than even the most primitive forms of representation. Thus, Burge suggests that there are non-psychological affective states. I think it is possible that some non-representational affective states are phenomenal even in creatures who lack representational psychology at all, since phenomenal consciousness may be more primitive than representation.

states are felt. Conscious perceptions of color, by contrast, are unfelt phenomenal states. A state's felt phenomenal quality may be localized (that is, associated with a spatial or proprioceptive representation), as when I have a pain in my leg, or it may be generalized, as when I feel tired.

There is a question whether emotions are always felt phenomenal states. If emotions are not just phenomenal, but also felt, then some states typically counted as emotions might not be, since they are qualitative, but not necessarily felt. Admiration may be one of these. Burge defines emotion as representational affect (2022: 546, fn. 771). This definition implies that emotion is not constitutively phenomenal, since neither representation nor affect, as Burge understands them, are constitutively phenomenal. It also implies that when emotions *are* phenomenally conscious, they are always felt, in my sense of felt, since affective states are always bodily ("visceral"). There are advantages to Burge's notion of emotion, not least of which is that it highlights the practical dimension of emotion, and helps make sense of the tight connection between emotion and attention.⁹³ Nonetheless, I prefer a notion of emotion that cleaves more closely to common usage, preserving the pretheoretical typing of emotions according to phenomenal quality, which suggests that emotions are constitutively phenomenal states. Thus, I hesitate to classify as emotion any state that is not, at least paradigmatically, phenomenal. Nonetheless, my critiques of Tappolet, including my arguments against the perceptual theory of emotion and representational neosentimentalism, do not depend on emotions being constitutively phenomenal. In fact, the difficulty of this issue provides further reason to resist giving emotion a central place in our value theory, and to focus instead on representation more generally and on phenomenal states.

⁹³ See Burge (2022, Chapter 15) and Brady (2013).

Finally, I think that phenomenal consciousness is the most fundamental kind of consciousness. This means that non-phenomenal conscious states are dependent upon a capacity for phenomenal consciousness, even though they do not have phenomenally qualitative elements as part of their nature.⁹⁴

V. Necessarily Phenomenal Modes of Presentation

I suggested above that some representations might have essentially phenomenal modes of presentation. Reflection on emotions and other affective states could warrant this supposition, which in turn leads to a conclusion about the ontology of value: Some values, aspects of value, or evaluative properties have natures that can only be represented through phenomenal modes of presentation.

What does it mean for a representation to have an essentially phenomenal mode of presentation? To answer this question, and to defend the claim that there are such representations, I begin with reflection on the attitude of admiration. These reflections will be especially helpful insofar as Tappolet also focuses on admiration as one of her central examples. Thus, what I say here will help clarify how my account is similar to and different from hers.

Admiration, as I understand it, is a kind of appreciation that typically takes rational agents and their qualities and actions as its proper objects.⁹⁵ I think that the most paradigmatic kind of admiration is admiration directed at a critically rational agent, like a human being.

⁹⁴ See Burge (2007a) and (2007b).

⁹⁵ We also admire things connected to agents and their excellent qualities. For example, one might admire a well-crafted table, or a painting. There may be other notions of admiration, as well, less obviously connected to agents, as when one admires beautiful scenery. These latter notions of admiration may be connected to the kind of admiration I discuss here, but for now I set them aside.

Similarly, the kinds of qualities that warrant admiration of an agent are those closely tied to that agent's rational agential properties. As Brewer (2009: 156) points out, it is bordering on incomprehensible to say that I admire a person for being two-legged. I think that this is because having two legs is not a property that bears directly on the essential agential functioning of the individual. That is, a person is the kind of agent that she is in virtue of having certain properties (capacities), and two-leggedness is not among them. The relevant properties picked out by admiration are constituted by various excellences of rational agency.⁹⁶

So far I have only sketched the kinds of things that admiration can be coherently directed toward, given its constitutive representational function. But even if this sketch is correct, I have not yet completely differentiated admiration from other kinds of representational state, since there may be other representational states that function to pick out the same class of objects—agents and their excellent properties, actions, and activities. As Brewer (2009) puts it, admiration “does not bring into view a generic sort of goodness but a special kind of value—call it merit.... [And] it is hard to see how admiration could possibly be reduced to the recognition of reasons for action [or appreciation]. Admiration involves a more direct, *essentially phenomenological* appreciation of merit or worth” (156–157, my emphasis). So, according to Brewer, part of the difference between admiring an agent and having the thought that the agent has merit is in the specific phenomenal quality of the appreciation involved. The plausibility of

⁹⁶ The nature of agency is teleological. It is constituted by natural norms. In the case of human beings, whose agency is representationally mediated, some of these are norms of veridicality: veridical perception, true belief, and the related actional representational notions of veridicality. These actional representational norms include norms for actional representations themselves—that is, representations, including motor commands and intentions, that are made veridical by the agent's activities—and norms for the correct application of attributives that function to be goals in actional representations—that is, evaluative representations. Admiration, since it is paradigmatically directed at agents and their essential properties (including their actions and activities), aims at a kind of value that is essentially teleological. Any other fitting objects of admiration must be, I think, parasitic on this central class of appropriate objects.

this claim provides some warrant for thinking that there are representations with essentially phenomenal modes of presentation.

Now, Brewer's concern in these passages is to reject what he calls the "buck-passing" account of value, according to which values can be reduced to reasons.⁹⁷ Reasons here are understood as non-evaluative properties that make various actions and attitudes appropriate. So on a buck-passing view, thinking about an agent *A* as having certain properties that call for admiration just is admiring *A*. But this misses the obvious distinction between *thinking* that *A* deserves admiration and *actually admiring A*. On my view, the feeling of admiration is a distinctive mode of presentation that is constitutively phenomenal.

The distinctive phenomenology of admiration is difficult to specify in abstraction. This difficulty is at the heart of the sentimentalist point that emotional experiences are epistemically essential to our understanding of value. If I am incapable of feeling admiration, then you will be hard-pressed to explain to me how that attitude presents its objects. Of course, you may describe to me the agent and her specific excellences qua rational agent, and I could thereby grasp the properties that constitute the agent's admirability. Nevertheless, there would be something about the agent and her excellence that I fail to grasp.

Consider the example of Socrates. I choose this example because I take it that my readers will find it familiar, though I know that Socrates may be somewhat controversial as an exemplar of admirability. For the sake of argument, I ask the reader to set aside questions of whether the Socrates presented in Plato's texts is a fair and accurate depiction of the historical person, and assume that he is indeed an admirable figure. I do not intend to take a stance on these matters here, but simply to take Plato's account at face value for the purposes of illustration. Of course,

⁹⁷ See Scanlon (1998) for an example of the "buck-passing" view that Brewer argues against.

the unsatisfied reader is welcome to conjure up another example from fiction, history, or personal experience. Socrates exemplified a kind of integrity, a consistent commitment to principles that manifests a special excellence of moral agency. He chose to die rather than act in a manner contrary to those principles. Moreover, the principles themselves—a clear-eyed loyalty to his country, honesty, the pursuit of understanding and the betterment of himself and his fellows—are worthy ideals around which to structure a life. With humor and wit, he faced his own end doing what he had spent the bulk of his life doing: humbly seeking wisdom.

Now, if I have done a good job of describing Socrates (or really, if Plato has), then you likely feel some admiration for him. And if you do not feel that admiration, then there is something about Socrates and his particular worth that you do not grasp. The simplest and strongest version of the idea is that the feeling of admiration is *the only way* to represent the relevant qualities of Socrates without loss, the only possible mode of presentation for fully and correctly representing his merit. A weaker version of the same idea claims that the only way to represent the relevant qualities of Socrates without loss is admiration *or another qualitatively similar mode of presentation*. It is somewhat speculative, but I shall take it as a working hypothesis, warranted by reflection on admiration and other similar attitudes, that at least the weak claim is correct.

Here an important question arises: Why is the phenomenal quality of the experience of admiration necessary for representing someone's admirability? What exactly does the phenomenal quality of the experience of admiration add to the content of the representation? One way to answer these questions is to suppose that the phenomenal mode of presentation is only necessary given our actual psychology. On this view, something about our psychology happens to work in such a way that evaluative representations can only play their functional roles by

virtue of having phenomenal modes of presentation. Another way to answer these questions supposes that phenomenal modes of presentation are necessary because of the nature of the value itself.

The first view, according to which evaluative presentations have phenomenal modes of presentation only as a matter of actual human psychology, has the advantage of not having to explain how an evaluative mode of presentation adds to the actual content of the representation. On this view, a phenomenal mode of presentation is necessary for representing an evaluative content only because that content always occurs in creatures like us in a phenomenal mode (or perhaps, to make room for the possibility of unconscious emotions and the like, in a mode that manifests in a phenomenal experience when attended to in the right way, or something like that). On this version of the view, the fact that some evaluative content occurs in a necessarily phenomenal mode does not tell us much about the nature of value, since the necessity involved is just psychological necessity.

On the other hand, phenomenal modes of presentation for some representations of value might be necessary not just psychologically, but ontologically. That is, it might be that *the nature of at least some kinds of value* necessitates a phenomenal mode of presentation. If this is right, then the phenomenal presentation of an evaluative representation adds something to the actual content of the representation. That is, a representation that applies the same evaluative attributive but in a non-phenomenal mode is either impossible or necessarily suffers from some kind of informational loss. The difficulty with this view is that its viability depends on the availability of an explanation for how exactly the nature of value puts this constraint on evaluative representation. And this explanation would need to draw on independently reasonable premises about both value *and* representational mind to clarify how a phenomenal mode of presentation is

the only way to capture a genuine evaluative property accurately. The success of such an explanation would serve as a strong argument in favor of the idea that some evaluative representations require phenomenal modes of presentation as a matter of ontological necessity.

So how might a phenomenal mode of presentation add to the content of a representation? Recall that phenomenology is essentially *qualitative*. Something is essentially qualitative when it has constitutive elements that cannot be fully explained in binary, mathematical, functional, or causal terms. An attributive is a purely formal, binary categorization, whose nature is constituted by the properties that it functions to be applied to. So if there are any attributives that can only occur in a phenomenal mode, this may be because the properties that they function to pick out have essentially qualitative elements. Suppose that a phenomenal attributive is applied to its object by virtue of having a mode of presentation that is *qualitatively similar* to the property that it functions to pick out. The qualitative features of reality, in other words, may be *mirrored* in the qualitative features of the mental presentations that constitute the representations of those features. And it is in virtue of this mirroring that the representations are able to represent those properties. Attributives are more than just open variables—arbitrary boxes, so to speak—into which particulars are sorted. Instead, the attributives are partly constituted by particular modes of presentation that connect the attributive to its referent through structural and qualitative similarity between the referent and the attributive's mode of presentation.

Burge (2022) calls representations that represent their subject matters through natural relations of similarity *iconic* representations, and he contrasts this iconic form of representation with purely conventional forms of representation, like those found in language. For example, a road map is an iconic representation of the actual roads and locations that it represents. A natural spatial and geometrical relationship exists between the lines and other markings on the map and

the things those lines and markings represent. By contrast, an English description of the same area represented by the map could also represent those roads and locations, but it would do so with representational elements that involve (largely) conventional relations between the representational elements (linguistic symbols) and the things those elements represent (the roads, locations, and the rest).⁹⁸ My proposal is that some representations have essentially phenomenal modes of presentation, and that the phenomenology of these representations bears an iconic relationship to the subject matter of the representations. That is, phenomenal modes of presentation at least sometimes represent their subject matters through natural relations of similarity between the phenomenal quality of the representational state and the subject matter being represented.⁹⁹

To see more clearly how some subject matters might necessitate particular kinds of presentation, and how reflection on those kinds of presentation can shed light on the nature of the subject matter, consider some examples from sense perception.

Visual perception, in order to be veridical, must involve modes of presentation with certain structural features that correspond to structural features of the physical environment. In visual perception, a two-dimensional array of information (the visual field) represents three-dimensional space. This, of course, is not the only way to represent three-dimensional space. Tactile perception represents three-dimensional space in a different way. But the structure of space itself dictates, to some extent, the nature of the subjective instantiations of mental

⁹⁸ See especially Burge (2022: 296–304). Notably, I am borrowing the term “iconic” and using it in a way that differs slightly from Burge’s usage. Burge builds the notion of *structural* correspondence into his technical definition of iconic representation (297–298). I am supposing that structural relations are not the only kind of natural correspondences that could hold between a representation and its subject matter. That said, Burge does allow for “very simple iconic representations” in which representational elements lack structure and nonetheless represent iconically (298). I am also not certain Burge and I understand the notion of “structural” in the same way, so there may be less difference between us than I am supposing.

⁹⁹ The notion of naturalness here is pretheoretical and intuitive. See Burge (2022: 300–301). Contra Burge, I resist glossing natural relations in terms of what is studied by “natural science” or mathematics.

content (the presentations) that can serve as representations of physical space. At the very least, the presentations must be rich enough to carry information about the size, distance, and so forth of particular objects.

Consider a more specific example of visual perception: color perception. Color experience is a phenomenal presentation that corresponds to electromagnetic wavelength information registered by the visual sensory apparatus. Any instantiation of mental content fit to represent variations in the electromagnetic spectrum must be sufficiently spectrum-like in order to carry the relevant information. It is conceivable, for instance, that some organism could represent these electromagnetic variations in a mode similar to human auditory perception, as a spectrum of tones, because that mode of presentation also has the necessary structure. Thus, we can understand color attributives as freely repeatable mental categories that apply to electromagnetic wavelength reflectance properties *by virtue of having modes of presentation with structural similarities to the electromagnetic spectrum*.

Now suppose that reality has purely qualitative elements in addition to quantitative ones. That is, suppose that some objects, properties, or relations are not identical with and cannot be reduced to mathematically specifiable structures, functional relations, or causal powers. Of course, some qualitative states also have mathematically structured elements. For example, pain states have intensity, color perceptions have hue, and so forth. And qualitative states can certainly have functions and causal powers. My claim is only that those mathematical structures, functional relations, or causal powers do not by themselves suffice to type-identify an essentially qualitative state. The nature of a qualitative state also contains non-quantifiable elements, such as the *feeling of pain*, that constitute the kind of state that it is.¹⁰⁰ My suggestion here is this: just as

¹⁰⁰ See Kim (2005), especially Chapter 6, for an argument that “qualia” resist functional or causal reduction. It is notable that this argument occurs in the context of a defense of physicalism.

some attributives are applied to their referents by virtue of having modes of presentation with iconic structural or mathematical similarities to their referents, some attributives could also be applied by virtue of having modes of presentation with *purely qualitative* natural similarities to their referents.

Some evaluative properties are essentially qualitative. So iconic attributives that function to pick out those evaluative properties will be applied via qualitative similarity, and thus will necessarily have qualitative modes of presentation. If this is on the right track, then the epistemic aspect of neosentimentalism is vindicated: emotions, understood as iconic evaluative representational states with phenomenal modes of presentation, are epistemically necessary. The point is that in order to represent without loss an essentially qualitative subject matter, one must have access to mental presentations with qualitative features similar to the qualitative features of the entities and properties that make up that subject matter. And thus, some representational states have necessarily phenomenal modes of presentation, and those modes of presentation are made necessary by the ontological nature of their evaluative subject matter.

This idea helps vindicate the emphasis placed by the evaluative outlook view on the phenomenal richness of certain evaluative representations. The “ineliminably evaluative representational content” of desire is something “to which we can return again and again in search of a deeper and more perspicuous understanding” (Brewer 2009: 154). Desires, as evaluative representations with phenomenal modes of presentation, could enable this deep cognitive access to the world of value if their phenomenal content iconically mirrors the evaluative properties of reality. And this feature of desire parallels the richness of perceptions,

which we can also return to over and over to improve our understanding of the subject matter perceived.¹⁰¹

Of course, I have only scratched the surface here, and left many important questions unanswered. The relation between phenomenal experience and motivation, the roles of attention and rational reflection in developing evaluative understanding, and the troublesome divide between goodness relative to function and inherent worth—all of these topics require treatment. In the remaining, I attempt to draw some of these threads together.

¹⁰¹ See Burge (2022: 360–363) for discussion of the iconic informational complexity of perception.

Chapter 4

Thought, Understanding, and Intention

The evaluative outlook model of practical reason claims that subrational capacities and mental states—such as emotion and desire—play a role in rational action. These capacities are subrational in the sense that they are more primitive than, and do not constitutively rely on, capacities for thought, inference, or critical reasoning. But in rational individuals, subrational capacities contribute to whole-individual functioning in a way that is nonetheless governed by rational capacities and states. It is constitutive of rational capacities that they function to play this organizing role in whole-individual functioning.

Dialectical activities illustrate this functional relationship between rational and subrational capacities. One is drawn into such an activity by an initially indistinct intimation of the activity's value, and engagement in the activity clarifies that value for the agent in a way that draws the agent into further stretches of engagement, which offers further clarification of the activity's value, and so on. The initial apprehensions of value that draw the agent into such an activity come from subrational processes. But these subrational processes are nonetheless shaped and guided by rational processes. The rational capacity that gives dialectical activity its dialectical structure is *understanding*. Understanding is constitutively rational because it relies on conceptual representational capacities. I think all central notions of understanding are closely related to knowledge, and so I count understanding as an epistemic good.

The species of understanding needed to account for dialectical activities must have several features. First, it must be able to take an activity and its constitutive ideals as an object.

According to Brewer, “each successive engagement [in a dialectical activity] yields a further stretch of understanding of the goods internal to the activity, hence of what would count as proper engagement in it” (2009: 37). Moreover, “one throws oneself into dialectical activities on the strength of an as-yet-indistinct intimation of their intrinsic value, hence without the benefit of a clear and definite representation of the mode of engagement that would properly consummate one’s desire” (37). Thus, the kind of understanding we are after must be able to manifest itself in inchoate representations of value. These representations will not be representations of states of affairs to be produced. They will not be capturable in terms of a non-evaluative description plus a separate evaluative valence (41). Rather, the evaluative element of these representations is an ineliminable aspect of the representational content. Finally, the representations through which this kind of understanding is manifested must be capable of moving an agent to act.

In what follows, I sketch an account of understanding that I believe satisfies all these conditions. The kind of understanding I have in mind is constituted by an individual’s ability to apply conceptual attributives to particulars. This is a general ability. It constitutes understanding of any concept that can be applied to particulars. I start by discussing the capacity of thought, which is necessary for any kind of understanding. Then I develop the notion of *basic understanding*, which I take to be a kind of understanding available to any organism with conceptual representational capacities. Next, I consider higher level thought, and how capacities for metacognition and critical reason enrich basic understanding. Finally, I turn to the practical dimensions of conceptual thought, proposing a notion of *intention* that coheres with the broader picture of representational mind I have presented in previous chapters. With this notion of intention in hand, I discuss rational motivation and attempt to dispel some persistent

philosophical confusions surrounding that concept. I conclude with reflections on how these distinctions can help us understand the nature of thought and its relation to the world of value.

I. Thought

Thought, as I understand it, is *conceptual* representation. Concepts are constituents of *propositionally structured* representations. Thus, a capacity for propositional representation entails a capacity for thought. Yet I think there is still an important distinction here, and so the two terms should not be used interchangeably. Propositional structure is *subject-predicate* structure. Concepts function constitutively to be elements in representations that have subject-predicate structure, but conceptual attributives can also be applied in more direct, perception-like representations. Not all conceptual representations have propositional structure, and so not all thoughts are propositional. Nonetheless, since a capacity for thought (conceptual representation) is dependent upon a capacity for propositionally structured representation, propositional representation is the more fundamental capacity. In this section, I develop an account of thought that respects these distinctions, with an eye towards applying that account to an understanding of practical reason.

Thoughts pick out particulars, properties, kinds, and relations (the objects of thought) by applying conceptual attributives in both singular (or plural) reference and in predication. (Perception picks out particulars in singular reference in the same way, albeit with a more limited array of attributives; the main difference is that *every* perceptual attributive functions to refer to particulars in context-bound singular reference, and in thought attributives need not always function to refer to particulars.)¹⁰² In thought, as in perception, attributives can be applied

¹⁰² See Burge (2010a: 380–381).

directly to the objects of thought in order to pick those objects out; in other words, reference is achieved by the application of attributives, which function to categorize particulars. For example, a thinker could pick out an object in thought as the brown dog. The thought refers to the object in part *by* attributing to it the property brown and the kind dog. This is not the same as referring to the object under a description like “the brown dog.” If it were, then reference would fail any time the description was inaccurate.¹⁰³ But reference can still succeed even when attribution, including attribution of essential properties, is not veridical. A thinker who refers in thought to an object as the brown dog still can succeed in referring to the object even if she is mistaken about its being a brown dog (though there must be *some* veridical element to her thought in order for reference to succeed—for example, that there is in fact an object). When reference occurs through description, the referent can only be whatever the description is true of. But when reference occurs through attribution, even though various properties are attributed to the object, the object is singled out independently from those properties. Nonetheless, the attribution is essential to the reference in the sense that the attribution is applied in order to achieve the reference (*it functions to refer*). It is impossible to refer to an object without making *some* kind of attribution.

Some representations that apply conceptual attributives have perception-like structure, applying the attributive in a context-bound way to a particular in the individual’s environment. For example, I might visually perceive an object in my environment *as* a brown dog. Such a representation need not have propositional structure; that is, I need not predicate anything of the object. I do *not* necessarily form a propositionally structured thought like that is a brown dog, even though brown and dog are conceptual attributives, which as such constitutively function to be applied in propositionally structured representations (representations with subject-predicate

¹⁰³ This is the same point made by Donellan (1966) and Kripke (1973), among others, about linguistic reference, but expanded to apply to reference in any representational state, not just in language.

structure). Because of how systems of propositional thought are integrated with sub-conceptual representation systems, like perceptual systems and intermodal non-conceptual representation systems, conceptual attributives might be applied directly by these other systems, or as a direct consequence of related attributions within these lower-level systems. (Of course, the details of whether and how this occurs, and in which specific representational systems, are empirical questions.) These lower-level representational systems do not generate representations with propositional structure. Instead of a representation with subject-predicate structure, I might form a representation with perception-like structure: brown dog there, or something like that. This sort of representation has the direct noun-phrase-like format characteristic of perceptions, without the predicate (or main-verb-phrase-like element) that marks propositional format. Such a representation might also import much of its mode of presentation directly from perceptual-level representations, especially if those perceptual-level representations apply learned attributives that are functionally connected to the relevant conceptual attributive. I count such perception-like representations as thoughts, and I call them *perceptual cognitions* (see Chapter 2 or this dissertation). A perceptual cognition like brown dog there might form the basis for a propositional thought with content like that brown dog is large.

Representations with perception-like structure, whereby particulars are picked out through the application of attributives, apply a kind of attribution that I call *intimate attribution*. Intimate attributions are those that apply attributives to particulars in order to pick those particulars out—to achieve reference.¹⁰⁴ In intimate attributions, reference and categorization go hand in hand. All perceptual attributions are intimate. Many attributions in propositional thought

¹⁰⁴ Intimate attributions are roughly *de re* attributions (see Burge 2009). But I avoid the label *de re* in part to minimize needless jargon and in part because I do not think intimate attributions are well-described by contrast with *de dicto* representation.

are also intimate. But propositionally structured representations also apply attributives *formally*, through predication. For example, in the propositional thought that brown dog is large, the attributives brown and dog are applied intimately, and the predicate is large applies the attributive large formally through predication. The application of the attributive large is not part of how the representation achieves reference to the particular in question. Constitutive or essential properties can also be applied both formally and intimately in thought. The propositional thought that dog is a terrier applies the concept terrier formally through predication to the object that is picked out via the intimate attribution of the concept dog. In the thought that dog is a terrier, the predicate indicates and refers to the kind terriers. That is, is a terrier indicates and refers to the kind terriers by invoking an attributive that functions to be applied to members of that kind, but not applying the attributive to any particulars. Of course, the proposition in its entirety does apply the attributive to a particular, but by itself the predicate is a distinct representational element that makes no such attribution to a particular. This ability to refer to kinds (and properties and relations) as such makes possible such thoughts as terriers are good dogs. In contrast, perceptual attributives do not function to refer to properties, kinds, or relations as such. They function to be applied to *particulars* in the environment that *instantiate* various properties, kinds, and relations. In this way, perceptions are always context-bound, while thoughts (like the thought terriers are good dogs) need not be.

The capacity to refer directly to properties, kinds, and relations—and not only the particulars that instantiate those properties, kinds, and relations—is one thing that makes conceptual representation different *in kind* from perceptual representation. This difference in kind is marked by a distance between the attributive and the object to which it is applied. Conceptual representation separates attribution from reference to particulars, making it possible

to arrange representations in subject-predicate format, and to refer to properties (and kinds and relations) directly. This separation is necessary for inference because inferences preserve truth relations in two ways: first, through the structure of propositions, and second, through preserving information about categories themselves. For example, from the thought that animal is a dog a thinker might infer that animal is not a cat. This inference could happen either because the thinker has a separate propositional thought, no dogs are cats, or because the thinker's concept dog somehow contains the information that cats are excluded from the category. For either kind of inference, the connection between the propositions is implicit in the representational structure and content of those thoughts. The inference to that animal is not a cat from that animal is a dog is contained in how those representations are presented in the mind of the subject. Importantly, this connection between premise and conclusion is not a matter of any explicit representational content; the thinker need not *represent* the one thought as following logically from the other. Rather, the two propositions are connected by a *structural* "because" relation in the subject's mind. In other words, a thinker can make an inference without having to (or even being able to) think about it *as an inference*.

Inferences, unlike perceptual transformations, are paradigmatically attributable to the individual. This is so even though not all inferences are consciously accessible. This marks another important difference between thought and perception. Both perceptions and thoughts are paradigmatically states of the whole individual, but the processes that produce perceptual states are not attributable to the individual, only to subsystems of the individual, such as the visual system (Burge 2010a: 540). By contrast, at least some of the processes that produce propositional states—namely, inferential processes—are also attributable to the whole individual.

There is nothing like inference at the level of perception. Relations between perceptual states are not contained in perceptual representations in any way: neither through explicit content nor through implicit structure. There may be attributives whose application excludes other attributives, as when an attribution of red excludes the simultaneous attribution of blue in color perception; but this is entirely a matter of the law-like processes that govern the formation of perceptual states, not at all a matter of the structure or content of those states. It may, however, be possible for perceptual cognitions (that is, representations that apply conceptual attributives but have perception-like structure—direct noun-phrase-like structure with no predicate) to be the basis of a proposition in an inference-like way. So, if a thinker has a perception-like representation with the content dog there, it might be possible for that thinker to form the propositional thought that is not a cat on the basis of that perceptual cognition. I am not sure whether this would, strictly speaking, be an inference. It does not involve the movement from one proposition to another, which I take to be the hallmark of inference. But it does involve information contained within or connected to the concept dog, and so leverages one of the central mechanisms of truth preservation. Perhaps that should lead to an amendment of how I have defined inference. Rather than insisting that inferences are always propositional representations formed on the basis of other propositional representations, I should say that inferences are propositional representations formed on the basis of conceptual representations (regardless of whether those conceptual representations have full propositional structure). I suspect, however, that the perception-like representation dog there must be transformed into a proposition like that is a dog in order to be the basis for an inference. Either way, these are empirical matters, and I do not think anything important in my arguments hinges on this point.

Possessing a conceptual attributive involves the ability to apply that attributive veridically to particulars, and some conceptual attributives, like tiger, apply to natural kinds as such. As Burge points out, possessing an attributive that applies to a natural kind as such implies “an openness to a distinction between surface properties and underlying natures” (2010a: 323). I think this “openness to a distinction between surface properties and underlying natures” is a constitutive feature of thought itself. Defending this claim requires additional reflection on the nature of thought. Recall that thought depends upon a capacity for propositionally structured representations with conceptual content. Propositional structure is subject-predicate structure; conceptual content is content that applies conceptual attributives, which are attributives that function in part to be applied as predicates. This characterization means that propositional thought is constitutively a capacity to separate attribution from reference to context-bound particulars.¹⁰⁵ A capacity to separate attribution from reference to particulars is a capacity to refer directly to properties, kinds, and relations. Any time a property, for example, is applied as a predicate, it is referred to directly and separately from any particulars to which it is applied. A thought with the content that body is square refers to the shape property squareness. In such a thought, the conceptual attributive square functions as a general and freely repeatable category, which can itself in principle be an object of thought.

Whenever a subject applies a conceptual attributive on the basis of perceived properties—as when a thinker sees an object as a piano on the basis of perceived size, shape, and color—the attribution is made on the basis of external, surface-level properties. Thus, there is a kind of surface level understanding that corresponds to a thinker’s fluency in applying conceptual attributives on the basis of such properties. For some attributives, these surface level

¹⁰⁵ See Burge (2010a: 539–547).

properties are also essential properties. For example, the conceptual attributive red refers to a property that is constituted by a certain perceptible appearance. And although there are interesting dependency relations involved in a full explanation of why red things are red, a thinker need not be able to comprehend such a full explanation in order to understand, in a basic sense, the property red. We pick out redness by whether something has a surface-level property of appearing red, and appearing red is also an essential property of redness.

But other conceptual attributives refer to kinds whose natures are not directly perceptible in the same way. The underlying nature of tigers is not directly perceptible. A thinker may apply the attributive tiger on the basis of perceptible qualities, like stripedness and growling, but applying the attributive tiger is not the same as applying an attributive like striped-growling-animal. The perceptible qualities are markers by virtue of which the thinker applies the attributive, but to apply the attributive is to identify the thing as having an underlying nature, *even if the attribution in and of itself has not connection any information about what that nature is*. A thinker does not need to know, or even have thought about, what features are constitutive of the kind tigers. The thinker need not be capable of identifying particulars according to the essential features by virtue of which the particulars are the kinds of things that they are. It nonetheless applies an attributive that functions to pick out that kind.

Now, some thinking subjects may have capacities for propositional thought that are limited in a variety of ways. It is worth reflecting on these potential limitations. It is not clear that all thinking subjects can predicate properties, kinds, and relations to the same extent. Perhaps some organisms have capacities for propositional thought that are not fully freed from the context-bound limitations of perception. Such a thinker might, for example, be able to think a partially context-bound thought like that is a cat, but be unable to form a fully generalized

thought like all cats eat mice.¹⁰⁶ But in any case, the gulf between a mere capacity for perception and a capacity for thought is wider than it may seem at first. Propositional thought involves not just a capacity to categorize features of reality, but also a capacity to think about those properties as distinct from the particulars that instantiate them. I think that this capacity—to separate attribution from reference to particulars—is the most fundamental element of thought, and it has surprisingly far-reaching implications. If an individual has a capacity for propositional thought, on this view, the individual can represent properties as such, without necessarily attributing those properties to particulars, thus enabling a slew of possibilities for mental content that would be impossible for individuals for whom attribution always functions to pick out particulars. For example, once properties are represented separately from their application to particulars, negation becomes possible. And negation makes possible question-structured representations.

Importantly, negation, quantifiers like some and all, and identity are not attributives.¹⁰⁷ They are structural elements of conceptual representation. In sufficiently sophisticated psychologies, they can themselves be represented, and so there can be attributives that function to pick them out. But conceptual representation systems need not have, nor be capable of having, attributives for these structural elements in order to generate representations that incorporate them. In a similar way, when an organism's representation system or action system applies an evaluative attributive, the organism need not be capable of applying attributives that have value or goodness as part of their content in order for the attributive to be evaluative in the structural and functional sense.

Suppose an organism with perceptual and conceptual representational capacities, but no capacities for metacognition or critical reasoning, forms a low-level representation as of a

¹⁰⁶ See Burge (2010a: 544).

¹⁰⁷ See Burge (2010a: 540–543).

physical body. Suppose further that it then forms a propositional representation with the content that is a body. This propositional representation contains an indexical, that, which picks out an environmental particular through perceptual reference. It also contains the concept body, which is a conceptual attributive directly derivative from the perceptual attributive for picking out physical bodies in the organism's immediate environment. But the capacity to apply the concept body, and not just the perceptual attributive, allows for a detachment of the organism's representational systems from the immediate objects that those systems function in this case to represent. For example, suppose that various features of the perceptual representation—perhaps the way the particular moves, or the way it interacts with light—cause the organism to withhold assent when it forms the propositional representation. The standard way to express this would be that the individual forms the thought that is a body without believing it. This withholding of assent is not possible in perception. The perceptual attributive body is either applied committally to the particular, or it is not applied at all. (There are complexities involving perceptual memory, anticipation, and imagination, but I am setting those aside for now.) Moreover, the organism might form the propositional representation that is not a body. This kind of representational form, negation, is not possible in perception (Burge 2010a: 540, 2022: 191). Again, the perceptual system either applies the attributive to the particular in a committal way, or it just does not apply the attributive. *But the organism need not be able to represent negation itself in order to form thoughts that have the propositional structure of negation.* Negation is, in the first instance, part of representational *form*, not content. The notion of “content” here refers to the attributives being applied in the representational state, as well as the mode of presentation and the veridicality conditions. Thus, there is no attributive not in the representation that is not a

body. Similarly, there is no attributive is in that representation. These are elements of the representational structure, not the content.

Propositional structure enables a capacity to form new kinds of non-committal representations.¹⁰⁸ Notice in the example described above that the ability to form non-committal conceptual representations allows the organism to be guided by something like curiosity. That is, the organism is guided by a representational goal of *confirmation*. There is surely a perceptual analog to the process of initial representation, exploration, and refinement. For example, a creature that merely perceives, but lacks a capacity for thought, might form an initial (nonveridical) perception with the content food there. It might then approach the object of its perception in an exploratory manner, and upon closer inspection, form a new representation of the object that does not apply the attributive food. But this process does not involve anything like negation. There is no perceptual content not food there. Moreover, the initial perception food there is fully committal, in the same way that a belief is fully committal. So there is no question-like structure. And insofar as the organism engages in exploratory behaviors, those behaviors are not guided by an initial representation that awaits confirmation; they are guided by a fully committal representation, which in this case ends up being dismissed. There is no sense in which the organism aims explicitly at confirming its initial perception. A conceptual representation, by contrast, can be formed noncommittally such that it can guide behavior that functionally aims at confirmation or disconfirmation.¹⁰⁹

¹⁰⁸ See Burge (2022: 190–201) for discussion of why the constitutively committal nature of perception precludes the possibility of perceptual negations, conditionals, or disjunctions.

¹⁰⁹ I make no claims about which organisms have these merely perceptual capacities and which have the more sophisticated capacities for curiosity. It may be that many organisms that seem to lack a capacity for thought do exhibit curiosity behaviors, which could suggest that their psychologies are more sophisticated than the label “merely perceiving” suggests. There may be gradations of complexity between mere perceptual capacities and full-blown propositional thought. Obviously, these are empirical questions. My goal here is to trace some of the boundaries between the various representational capacities, not to issue assertions about the actual capacities of particular organisms and species.

These abilities are more important than they seem. A subject with these capacities can form representations that require confirmation, which may slow down and complicate the process of action. Of course, perceptual capacities also complicate the agential process by mediating between stimulus and behavior, and making possible actions that are evaluated in terms of veridicality. But propositional thought creates space for *curiosity*, and marks the beginning of a sensitivity to a seems/is distinction. I do *not* mean that thinking subjects must have a capacity to *represent* a seems/is distinction; rather, a capacity to make non-committal conceptual attributions is a necessary precondition for representing a distinction between appearances and reality. This capacity for non-committal attributions marks the lower border of an ability to discern surface appearances from underlying natures.

II. Basic Understanding

Propositional thought represents a substantial advancement in psychological sophistication over and above the capacity for perception (though perceptual capacities are often very sophisticated, and likely mark the lower border of having a psychology at all). But human beings have additional capacities even more sophisticated than propositional thought. Human beings (and perhaps some other animals) have capacities for *understanding* and for *critical reasoning*. In this section and the next, I explore these additional capacities and fit them into the larger picture of representational mind that I have been sketching so far.

I begin with understanding. First, I do not think that understanding necessarily is a more sophisticated capacity than thought. In fact, I think that one kind of understanding comes online as soon as an organism becomes capable of conceptual representation. But other notions of

understanding do seem to be more sophisticated, requiring either articulacy (the ability to give and understand explanations), or certain kinds of conceptual content (like content that includes other concepts, propositions, or whole networks of inferentially connected propositions). Let us start with the less sophisticated kinds of understanding, since I think the more complex sorts of understanding build on those more basic capacities.

I begin with what I take to be an element of common sense: A thinker's ability to refer to an object in thought under various attributes depends on the thinker's understanding of the object and those attributes. For example, a thinker unfamiliar with dogs might think of a dog in her environment as that animal, but be unable to think of it as that dog, either because she lacks the attributive dog, or because her ability to apply the attributive dog to novel objects is limited. Such a thinker might become able to form thoughts that accurately predicate the attributive dog to the animal. Suppose, for instance, someone tells her, "That is a dog," and she thereby thinks that animal is a dog. But notice the difference between this and the ability to think that dog. The latter is a direct application of the attributive dog, and the former is not. Similarly, a young child might think of all furry four-legged animals as dogs prior to mastering a more ample array of attributives. Someone could correct the child in a particular case and say "That is a cat." The child might then think that is a cat, but still think of the next cat he encounters as a dog. The child might even think that dog is a cat. These are clearly cases of failed or underdeveloped *understanding*.

Let us call this *basic understanding*: the ability to apply a conceptual attributive fluently. "Fluency" here is akin to linguistic fluency. It includes not only the ability to apply a conceptual attributive veridically to particulars, as when a speaker uses the right word in the right context, but the ability to apply that attributive in novel contexts and to refrain from applying it when it is

not appropriate. Moreover, fluent application of a conceptual attributive is relatively effortless on the part of the thinker. The thinker who has basic understanding of a concept does not need to deliberate consciously in order to apply the conceptual attributive veridically to most particulars. Rather, the attributive has been integrated into the thinker's representational systems so that it can be applied directly on the basis of relatively low level representations (like perceptions). But the objects of basic understanding need not be concrete particulars. Paradigmatically, they include abstract objects and properties, like numbers, kinds, or relations. To have basic understanding of the kind dogs, for example, is to be able fluently to pick out dogs in perceptual reference and form true propositional thoughts about dogs directly on the basis of perceptual representations.

Basic understanding manifests itself in both perception-like attributions—that is, context-bound attributions to particulars in the individual's immediate environment—and pure attributions—that is, attributions that occur separately from any one context or particular.¹¹⁰ Both these kinds of attribution are elements of thought. Let us consider some examples.

Suppose I form a perception-like representation, dog there. Such a representation applies the conceptual attributive dog intimately; that is, the attributive is applied directly to the object, using a noun-phrase-like structure, in order to achieve reference. The conceptual attribution occurs directly on the basis of lower-level perceptual attributions, perhaps involving learned perceptual attributives such as dog-shaped or dog-body. The conceptual representation therefore inherits a perceptual mode of presentation. Because dog is a conceptual attributive that is applied directly on the basis of perceptual attributions, this representation is what I have called a perceptual cognition; it has perception-like structure, and may be generated in a relatively

¹¹⁰ I borrow the term “pure attribution” from Burge (2010a). In later writings (2022), Burge simply uses the term “predication” and speaks of “main-verb-phrase-like” structures instead of pure attributions.

low-level representational system, but it applies a conceptual attributive, an application which depends on a capacity for propositional thought. This representation is a manifestation of my basic understanding of dogs. A subject with a capacity for thought can also form representations that are not context bound. I can think about dogs in general, for example, or I can contemplate the dog I saw three days ago. Such thoughts also employ the conceptual attributive dog in order to achieve reference, so they employ intimate attributions, but this time the attributive is not applied to a particular in the individual's immediate environment. Thinking about dogs in general uses the attributive in order to refer to the *kind* dogs. Thinking about the dog I saw three days ago applies the attributive in order to refer to the specific animal I saw at that time. Notice that in the first case, the attributive is not applied to particulars at all; no particular object is thought of as a dog. In the second case, the conceptual attributive is applied to a particular, but that particular is not currently present in the thinker's physical environment. Moreover, unlike with a perceptual cognition, the application is not directly based on currently applied lower-level perceptual attributives (though perceptual memory is almost certainly involved). Both these kinds of employment of conceptual attributives are examples of non-context-bound attribution.¹¹¹ The first case is also an example of *pure attribution*, in that the attributive refers to a property, kind, or relation without applying to any particulars that instantiate that property, kind, or relation. It is a hallmark of conceptual attributives that they can be employed in pure attributions. And any propositionally structured thought—such as that is a dog—involves pure attribution since the predicate, or main-verb-phrase-like element of the thought, attributes a property but has no

¹¹¹ Merely perceiving organisms can refer to objects not currently present in their immediate physical environment through perceptual memory, much like my thinking about a dog I saw yesterday. But of course, their doing so would not involve a conceptual attributive, and so they would not be capable of pure attribution. For extensive discussion of perceptual memory, see Burge (2022: chapters 16 and 17).

referential function. Thus, is a dog applies the attributive dog to the subject-place object, but has not essential role in achieving reference to that particular.

Fluency in the application of conceptual attributives constitutes a kind of understanding. Understanding is often thought of as an *ability*.¹¹² My account here follows suit. My ability to apply the attributive dog correctly in various representations partly constitutes my understanding of dogs. Thus, understanding is a psychological capacity or competency. It constitutively involves application of conceptual attributives. Organisms that lack capacities for conceptual representation are not capable of any kind of understanding. The objects of basic understanding are also the objects of thought: particulars, properties, kinds, relations. In general, basic understanding does *not* apprehend logico-linguistic relations or relations among representations.¹¹³ And even though there is a sense in which basic understanding is understanding *of concepts*, concepts as representations are not the objects of basic understanding. Rather, the particulars, properties, kinds, and relations that concepts function to represent are the objects of basic understanding. Insofar as logical relations are aspects of basic understanding, they are so because the conceptual representations that constitute basic understanding instantiate those relations when they function as elements in propositionally structured thoughts, not because those representations take logical relations as part of their content.

Basic understanding can also manifest in fluent *inferences*, which are logical (truth-preserving) transformations between propositional representational states. Inferences preserve truth by exploiting either the structural elements of propositions or the information

¹¹² See, for example, Elgin (2007) and Hills (2009).

¹¹³ See Kim (1994) for an exploration of internalist accounts of understanding, according to which understanding fundamentally takes logico-linguistic relations as its object. I think basic understanding, as I have described it, constitutes a fundamental kind of understanding, and so I reject such internalist views. There are, however, special cases, such as understanding of mathematical, logical, or linguistic subject matters.

attached to concepts. One kind of general inferential fluency involves manipulating propositional structure. Such structural inferences are independent of conceptual content. For example, the inference from All Fs are Gs and E is an F to E is a G is valid independently of the concepts that fill in the variables. But inferences can also be based on relations between concepts, and these inferences do rely on mastery of conceptual content. It is this kind of inference that manifests basic understanding. For example, the ability to infer from the thought that is a dog the further thought that is not a cat is a manifestation of basic understanding of the kind *dog* (and the kind *cat*). Moreover, the ability to make this inference fluently, without conscious effort or deliberation, shows greater basic understanding than if the same inference were to require conscious effort.

Now one might reasonably suppose the inference from that is a dog to that is not a cat demonstrates the thinker *knows* that no dogs are cats. In a loose sense, I think this is correct. But I doubt such an inference necessarily depends upon the token propositional thought no dogs are cats occurring anywhere in the individual's psychology. I am not even sure the thinker must be capable of forming the thought no dogs are cats in order to make the inference. Logical quantifiers like no, some, and all are not constitutive of a capacity for propositional or conceptual form as such. It is conceivable that some primitive psychologies lack such quantifiers, and nonetheless include a capacity for basic perceptual beliefs and inferences. Such a thinker would be capable of forming thoughts like that is a dog based on perceptual representations translated into conceptual format, but would be incapable of forming any thoughts about all dogs. Nonetheless, I suspect that such a thinker would be capable of an inference like the one from that is a dog to that is not a cat—assuming, of course, the individual has a capacity for conceptual negation. But a similar argument can be made about conceptual negation, too. If the individual

lacked that capacity as well, still inferences manifesting basic understanding would be possible. For instance, the individual could in principle infer something like that is dangerous from that is a dog, even if the individual lacked access to *any* logical connectives. Thoughts like that is not a cat or that is dangerous can be *justified* directly by a prior thought like that is a dog. The formation of one of the former thoughts on the basis of the latter is a paradigmatic instance of justification because the former thoughts are reflectively accessible grounds of the latter.¹¹⁴ The premise, the conclusion, and the inferential transformation are all attributable to the individual. I conclude from these reflections that conceptual content can be the engine of inference.

I am tempted by this line of reasoning to say that understanding is not a species of knowledge.¹¹⁵ Knowing that *p* seems to imply at least a capacity to form a thought with the content *p*, even if such knowledge is supposed to be tacit or implicit. Basic understanding, as I have characterized it, carries no such an implication, in large part because it does not take *propositions* as its objects. (Of course, there are characterizations of knowledge that allow objects other than propositions, as when one knows a person or knows a place. Depending on whether such locutions indicate genuine types of knowledge, basic understanding may yet be explainable as a species of knowledge.) Because basic understanding consists of capacities centering around conceptual content, I tend to view *conception* and *understanding* as roughly the same thing. Thus, I do not think a conception should be understood primarily as a network of propositions surrounding a concept, at least not if the constituents of such a network must be thought or thinkable by the individual. Nothing in my larger argument hinges crucially on these points, however.¹¹⁶

¹¹⁴ I follow Burge (1995, 2003) in this characterization of justification.

¹¹⁵ See Grimm (2006) for an argument that understanding is a kind of knowledge.

¹¹⁶ My notion of basic understanding, and my arguments in these preceding two paragraphs, may be related to Higgenbotham's (1998) distinction between concept, conception, and conscious conception. Indeed, one might

Basic understanding is a representational competency based on the capacity for conceptual attribution. Conceptual attribution answers to the particulars, properties, relations, and kinds that actually exist in the world. This point falls out of the general point that representations function constitutively to be veridical. A conceptual attribution will not be veridical if it is applied too widely or too narrowly. Consider the child who applies the concept dog to cats; similarly the child might fail to apply the concept dog to a chihuahua. In the first case, the concept is applied too widely, and in the second, too narrowly. This example, although illustrative, is also somewhat misleading, because the human capacity for language makes it easy to confuse concept development with language acquisition. The two are closely related, but not identical.¹¹⁷ Moreover, dog is a somewhat difficult concept to use as an example since dogs are not an unambiguous example of a natural kind. So let us develop another example.

Suppose an ape has a capacity for conceptual thought but has no language. The ape has a concept tiger. The ape does not use any arbitrary linguistic symbol to mediate the application of this concept—that is, the ape has no word “tiger” and no linguistic community to determine in advance the application conditions for the concept (that is, the concept is not associated with vocal calls, gestures, or other arbitrary perceptible symbols). Suppose the ape applies the concept tiger to a housecat. This is a nonveridical application of the concept—it applies the concept too

reasonably think that my notion of basic understanding is very close to Higgenbotham’s idea of conceptual competence. His discussion of conceptual competence and implicit conceptions, in which he considers how reflection on a concept can reveal truths that one “knew all along” but never articulated to oneself, seem broadly congenial to my central claims about basic understanding and conceptual content here. Higgenbotham also presents cases of employing a concept inadequately or with significant misconceptions, which resemble the phenomena I present here as instances of misunderstanding. I would resist, however, his glossing of conceptual competence in terms of “know[ing] the nature of [one’s] own concepts” since that implies an articulacy about one’s conceptions that need not be present for one to possess basic understanding. I must admit that I lack a good grasp of all the relevant distinctions in this vicinity, and I think further reflection on these matters would be fruitful.

¹¹⁷ I think that capacities for thought always precede capacities for language. Though language certainly enhances capacities for thought in some ways, the basic representational capacities must already be present in order for language to work the way it does.

widely. Such an application, I suggest, is a failure on the part of the ape to understand the *kind*, tigers, that the concept constitutively functions to pick out. Similarly, if the ape fails to apply its concept tiger to a white tiger, and perhaps instead employs a new concept tiger* to pick out the white tiger, this too would be a misunderstanding, since the white tiger is in fact of the same kind as the other animals which the ape's concept tiger functions to pick out. This latter case would be an example of applying the concept too narrowly.

The notion of basic understanding has the advantage of explaining one aspect of *depth* of understanding. Depth of understanding, on this view, can be accounted for in terms of how well a conceptual attributive has been integrated into the representational psychology of the individual. This integration is a matter of how fluently the individual's representation systems—including non-propositional representation systems—apply the concept. If a conceptual attributive is well-integrated with the individual's overall representational psychology, the individual will generally apply the attributive to particulars that in fact fall under the concept and not to particulars that do not fall under the concept. Moreover, some of these applications occur directly on the basis of sub-propositional representations; they are not predications, and they are not inferences. Rather, they are perception-like intimate applications of the conceptual attributive. Intimate applications function to pick out particulars through attribution, and perception-like applications typically occur through processes that are not attributable to or reflectively accessible by the individual, though the resultant representational states, the outputs of those processes, *are* typically attributable to the individual. Finally, these applications of the conceptual attributive do not require deliberation or conscious effort on the part of the individual. These are the marks of conceptual fluency. And depth of (basic) understanding is marked by a high degree of fluency with the concept in question.

When the ape encounters a tiger, suppose that the ape forms a representation with the content tiger there. This content has perception-like representational structure. It is not propositional. It achieves reference to the particular (the tiger) through perceptual processes that result in the application of the conceptual attributive tiger. The application of that conceptual attributive is based directly on lower-level visual attributives, which pick out properties like shape, color, size, and motion. Those same lower-level attributives might also be the basis of other higher-level attributions, such as body, animal, and even tiger-body.¹¹⁸ Some of these attributives, like body, occur within the visual system itself; others might occur in other representation systems and are dependent upon propositional representational capacities. Regardless of the exact psychological details, the representation tiger there occurs automatically, without conscious effort on the ape's part. It is not an inference—inferences result in propositionally structured representations, and this state is not propositionally structured. The processes involved in the formation of this state are not reflectively accessible to the ape, and the processes themselves are not something that the ape *does*. Instead, those processes occur within the ape's representational subsystems. Nonetheless, the output of those processes, the representational state, *is* attributable to the ape.

I think the ape's representational state is a manifestation of the ape's basic understanding of tigers. Notice that the object of this understanding is the *kind*, tigers, not the concept itself, or some other internal feature of the ape's psychology; nor is the object anything like logical relations between propositions or concepts. It would be highly implausible to suppose that the ape must be able to represent the concept *as such* in order to understand the kind to which the

¹¹⁸ The locution tiger-body indicates a perceptual attributive that functions to pick out bodies of a certain shape, size, color, and so forth. It is not meant to imply that perception can pick out the natural kind *tigers* as such, though perceptual learning may sometimes enable such natural kind sorting. See Burge (2022: 118–130, 499–501). Burge also describes a general taxonomy of levels of perceptual sorting (2022: 427–431).

concept functions to refer. Similarly, it would be implausible to suppose that the ape must be able to represent logical relations between propositions or concepts. Pretheoretically, the ape's ability to do things like correctly identify tigers and make inferences about individual tigers on the basis of their membership in that category need not rely on such high-level representational capacities. For example, if the ape has the ability to think a thought like that is a tiger, and to infer from such a thought further thoughts like that is not a housecat and that is dangerous, these abilities do not depend on the ape's having any metarepresentational abilities to think about the concept tiger as such, or to think about propositions, inferences, or reasons as such. And these more basic conceptual, propositional, and inferential abilities can partly constitute the ape's understanding of tigers, even in the absence of such metarepresentational abilities.

My point here is that picking out particulars as tigers in representational states (both propositional and non-propositional states), including the application of the conceptual attributive tiger directly on the basis of lower-level perceptual attributives, also constitutes basic understanding. But understanding the concept tiger, in the sense that I mean, is about understanding the *kind* that the concept functions to refer to. Thus, to use Kim's (1994) terminology, the notion of understanding that I am developing here is thoroughly *externalist*. That is, the objects of understanding are aspects of the world, *not* features of the subject's psychology or epistemic state.

Basic understanding helps explain how some of our representational states can seem at once to be our own and not our own (Tumulty 2020, especially chapter 1 and 2). The concepts that make up our basic understanding are incorporated into inferential networks and are generally sensitive to rational reflection. Most of the representational states in which they figure are attributable to us as individuals, and in that sense they are our own. But the processes that

produce many token applications of those concepts are directly based on perception-level representations, and thus are not attributable to us as individuals, subject to volition, or accessible to reflection; so they feel like they are not our own. These points go some way toward explaining away the so-called “problem of receptivity” (see Chapter 2 of this dissertation, especially *Section IV: Desire and Problem of Receptivity*). Basic understanding also helps explain the dialectical structure of many human activities, since that structure depends upon an interplay between understanding and perception-like representations of value.

I do not claim, however, to have presented a comprehensive account of understanding. Rather, at this point I have only developed a specific notion of understanding, defined in terms of particular representational abilities—namely, abilities to apply concepts and conceptual attributives fluently. But there are more sophisticated types of understanding.

III. Metacognition, Critical Reason, and Theoretical Understanding

Some thinkers achieve a level of understanding that clearly goes beyond the basic understanding that I described in the last section. This higher-level sort of understanding is often about whole subject matters, and it involves possession of a *theoretical framework* for thinking about the subject matter. While it is clear that basic understanding is something achievable by young children and many non-human animals (the ones with a capacity for propositional thought), this *theoretical understanding* may be achievable only by thinkers with more sophisticated representational capacities.¹¹⁹

¹¹⁹ See Grimm (2016) for an argument that understanding must be something achievable by young children, and probably some non-human animals as well. As we shall see, there may be types of what I am calling theoretical understanding that are available to such thinkers as well.

Consider this passage from Elgin (2007):

To understand [a complex event or subject matter] involves more than knowing the various truths that belong to a suitably tethered comprehensive, coherent account of the matter. The understander must also grasp how the various truths relate to each other and to other elements of the account. She should also be able (and perhaps be aware that she is able) to use that information, to reason with it, to apply it, perhaps to use it as a source of working hypotheses about related matters. Someone who knows geometry, for example, knows all the axioms, all the major theorems, and how to derive the major theorems from the axioms. You can acquire this knowledge simply by memorizing. But someone who understands geometry can reason geometrically about new problems, apply geometrical insights in different areas, assess the limits of geometrical reasoning for the task at hand, and so forth.
(35)

If theoretical understanding requires things like “assess[ing] the limits of geometrical reasoning for the task at hand” and an awareness of one’s own abilities to use information, then it seems to require not just a capacity for thought, but a capacity for *critical reason*.

Critical reason is a capacity to represent mental states as such—beliefs, intentions, inferences, and so on—and also to evaluate them according to the constitutive norms that guide them. Reasoning, as I have described it, is simply the ability to make inferences: to form new propositionally structured representations on the basis of other conceptual representations.¹²⁰ Anytime a thinker forms a belief or an intention on the basis of a conceptual representation, that thinker is responding to a reason. Critical reasoning is a capacity to assess how successful a belief or intention is according to the (rational) norms that constitutively function to guide it.

¹²⁰ In the next section I discuss the practical dimensions of reasoning. But suffice it to say for now that actions and intentions can be just as much the products of inferences as beliefs can be.

Thus, reasoning can occur without any capacity for metarepresentational abilities; critical reasoning cannot. An example will be helpful.

Suppose a mouse possesses a capacity for propositional thought, but no metacognitive representations at all. The mouse will be capable of making simple inferences, as when it thinks that animal is a cat and infers that animal is dangerous. The mouse's second belief follows from the first, and the first provides warrant (in this case, justification) for the second. Neither of these facts depend on the mouse *conceiving of* the first belief as a reason for the second. Nonetheless, the first belief acts as a reason for the second in the mouse's psychology. The mouse is employing a capacity to reason. By contrast, a human being who has a similar pair of beliefs, that animal is a tiger and that animal is dangerous, can also represent the first belief explicitly as the reason for the second. Moreover, the human being can assess whether the first belief is a *good* reason for the second. Thus, the human being not only has a capacity to reason, but a capacity to reason critically.

Before returning to my discussion of understanding, I will pause briefly to consider the relation between metacognition and critical reason. Obviously, metacognition—the capacity to represent beliefs, intentions, reasons, or other mental states as such—is necessary for critical reason—the capacity to evaluate mental states according to their constitutive functions. It is conceivable that some thinkers might possess metacognitive capacities without critical reasoning capacities. Nonetheless, the connection between metacognition and critical reason is a tight one. It is hard to imagine, for instance, possessing a concept of belief without also possessing a concept of truth, the constitutive norm of beliefs. This is simply because it is hard to imagine a thinker being able to distinguish beliefs from the states of affairs that those beliefs represent without some notion of a seems/is distinction, and thus some notion of correspondence between

representation and reality. The same reasoning seems to apply to any representational notion. Thus, it is tempting to suppose that any metacognitive capacity comes with a corresponding ability to distinguish between successful and unsuccessful representational states, at least in a rudimentary way.

On the other hand, there are psychological states that are not constitutively functional. Conscious states like pain have no constitutive teleological structure. A psychological state is a pain state by virtue of its phenomenal quality, and it need have no representational content, nor serve any functional purpose whatsoever, in order to be a pain state. Of course, pain states do typically serve some functional purpose, such as biological fitness, but such functions are not constitutively part of the nature of the pain state in the way that the truth condition is constitutively part of the nature of a belief state. Phenomenal consciousness as such is not constitutively functional. Thus, states whose natures are entirely determined by conscious phenomenal quality are not in themselves functional. And so there is no such thing as a pain state that succeeds or fails *as a pain state*. Any notion of success or failure that applies to a pain state must be imported from some external functional system in which the pain state plays a role.

Perhaps, then, there are organisms who lack any concepts for representations as such but who nonetheless have some psychological concepts. Suppose a dog, for example, has a concept of pain. The dog may even be able to apply that concept to pain felt by other animals. Should we call this a capacity for metacognition? It is certainly a capacity for meta-*psychology*, and it would constitute a significant layer of sophistication in an animal's representational abilities. I think that the notion of metacognition, however, most appropriately applies to representations of *representational* states. And although metacognition and critical reason may come apart, and there almost certainly can be differing degrees of metacognitive and critical reasoning abilities, I

will proceed under the presumption that the most significant distinction in the vicinity is that between mere propositional thought and critical reason, without being overly concerned with the precise location of the border between these two capacities.

So understanding of the kind that Elgin (2007) describes, which I call theoretical understanding, requires critical reasoning capacities. This requirement corresponds to a requirement that the one who possesses theoretical understanding has cognitive access to, and an ability to evaluate, *explanations*. An explanation, as I use the term, is a representation, or cluster of representations, that contains a “because” relation as part of its content. This “because” relation is the relation of *rational causation*, understood as an inferential connection between one conceptual representation and another.¹²¹ A thinker who can explain some proposition *p*, not only believes that *p* because she believes that *q*; she also believes that *p because q*, which involves representing that *q* provides a reason to believe *p*.¹²²

There are two kinds of rational causation at play here. On the one hand, a thinker might see *q* as a reason to believe *p* contingently, as when I believe that you are not at home because I rang your doorbell and received no answer. In this kind of case, *q* (I rang your doorbell and received no answer) provides warrant for *p* (you are not at home), but it only explains *p* in the sense that it explains why *I believe* that *p*. There is no necessary link from the fact that you did not answer your door to the fact that you are not at home, even though the former constitutes evidence for the latter. If, on the other hand, I believe that you are not at home (*p*) because you are at the post office (*q*), my belief that *q* explains my belief that *p* not only in the sense that it explains why *I believe* that *p*, but also in the sense that it explains why *p is true*. The fact that you

¹²¹ See Marcus (2012) for an extended discussion of the notion of rational causation.

¹²² Burge describes reasoning as “a propositional sequence that functions to provide an explanatory justification of a conclusion” (2022: 527). But it is not clear to me what exactly he means by “explanatory justification” here. My discussion below aims to clarify some different kinds of explanation.

are at the post office *entails* that you are not at home. I shall call the first kind of explanation, in which q explains only why the thinker believes p , *contingent rational explanation*; I shall call the second kind of explanation, in which q also explains (at least partially) why p is true, *necessary rational explanation*.

This distinction complicates the claim that theoretical understanding requires critical reason, or even that it requires metacognition. Certainly, representing contingent rational causation requires at least metacognition (and if that representation also involves an assessment of success, then also critical reason), since contingent explanations are explanations of beliefs (and other propositional representational states). But necessary explanations may not require metacognition at all. Of course, if necessary explanations include contingent explanations, then they also require metacognition. As I presented them, necessary explanations explain *both* why a thinker believes that p and why p is true. And insofar as an explanation of why p is true also provides a reason to believe that p , that is correct. But it is possible that a thinker could have no ability to represent mental states as such (that is, the thinker would possess no conceptual attributives that function to pick out mental states, so no concepts of perception, belief, intention, reason, and so forth), and yet be able to represent causal or dependency relations.¹²³ Thus, it seems quite likely that some thinkers can represent a kind of “because” relation without any metacognitive representational capacities. If that is right, then some thinkers can cognitively apprehend the key aspect of necessary explanations, and they can do so without metacognition, much less critical reason.

I suspect that this representation of a “because” relation is sometimes a matter of the logical form of the representation, rather than a matter of a conceptual attributive that picks out

¹²³ Burge thinks it is likely that mechanical causation can be represented even at the perceptual level (2022: 446–466). See also Grimm (2016).

causal or dependency relations. Thus, a representation with content like *p* because *q* is akin to a representation of not *p* in that the “because” and the “not” elements are structural features of the representation, and need not involve any attributives that pick out those relations. But even if an organism does have attributives for picking out causal or dependency relations, that does not itself entail metacognition. Causation and dependence are not representational or psychological notions. Thus, it seems possible that a thinker can understand explanations without any metacognitive capacities. And so explanatory understanding, which involves grasping of explanatory relations, probably comes in metacognitive and non-metacognitive flavors. I propose reserving the term “theoretical” for the metacognitive kind. Thus, we can say that whenever a thinker correctly represents dependency relations as part of its conception of an object (where “object” is broadly construed to include particulars, kinds, relations, events, and subject matters), that thinker possesses some degree of *explanatory understanding* of that object. If the thinker’s conception also contains representations of the conceptual contents as such, perhaps along with assessments of the adequacy of various conceptual elements that make up that conception, then the thinker possesses *theoretical explanatory understanding*.

I am not sure how much my overall account of evaluative outlooks hangs on these distinctions. But I think they provide useful background for the explanations of practical thought and motivation that follow.

IV. Intention and Rational Motivation

Having sketched a framework for thinking about propositional thought, including several notions of understanding, I turn now to the question of how these various capacities for

higher-level thought bear on agency. I argue in this final section that a capacity for propositional thought mediates agency by creating a governing level of conative states that function to organize lower-level conative states.

Recall that some attributives, like food and mate, function to serve as goals in action commands. And action commands, in their most primitive form, function to guide bodily movements by mapping out those movements and employing behavioral schemas to execute the movement so mapped. Conative states include both these goal setting attributives, which are evaluations or proto-evaluations, and action commands.

Functional attributives have veridicality conditions of the same kind as other attributives. The attributive food is veridical if the particular it applies to is in fact food for the organism, just as the attributive body is veridical if the particular it applies to is in fact a physical body.¹²⁴ Action commands have a different kind of correctness condition. Action commands are veridical insofar as the bodily movements that they map are executed in accordance with the representation.

A capacity for propositional thought enables a different kind of action command, one that is not context-bound. Rather than merely mapping out specific bodily movements or behavioral schemas, conceptual action commands can map out actions according to general categories of action without tying those action-types to an immediate context. For example, a spider that lacks conceptual representational capacities, but has perceptual capacities, forms action commands that map out particular instances of hunting behavior, but the spider cannot represent hunting as such, decoupled from its current hunting activities in the here and now. That is, all of the spider's representations of hunting are action commands for directly guiding particular instances of

¹²⁴ Of course, the functional attributive has veridicality conditions that are relative to the nature of the organism whose representation it is.

hunting behavior. Suppose that a wolf has a capacity for conceptual representation. The wolf forms action commands for particular bodily movements that constitute hunting behavior, just like the spider, but it also might form conceptual action commands to hunt. These conceptual action commands guide the wolf's behavior, but not by directly mapping specific bodily movements and behavior schemas in the immediate context. Rather, these conceptual action commands govern which other action commands (the ones that do map specific context-bound bodily movements) are formed. This controlling function of some action commands may be achieved through a variety of means: focusing perceptual attention for the purpose of goal-setting, directly causing lower-level action commands, prioritizing certain kinds of lower-level action commands over others, and by preventing some lower-level action commands from forming. (These are, of course, empirical matters. Here I merely speculate about possible mechanisms.) This kind of higher-level action command is an *intention*. Because they are more removed from the lowest level motor commands, intentions give agents more control over their actions.

To be clear, perception level action commands also contain levels of generality. Some perceptual-level action commands take something like infinitival form, such as to hunt, which govern more basic motor commands for particular bodily movements. The distinction between perceptual-level action commands and intentions is a distinction between different psychological kinds, differentiated by different forms and functions. But there are analogs between the two kinds of conative state. Burge (2022) puts the point as follows:

There is [a perceptual-level] analog of a propositional-level intentional state, or other propositional goal-setting conative state. There is a further motor command that is [a perceptual-level] analog of an occurrent implementation of a propositional-level intention. Such states depend on specificational states that function

to provide information about conditions under which the act is to be carried out. These latter states (perceptions, perceptual memories, perceptual anticipations) are analogs of beliefs. (527)

So an intention is a conceptual action command. Intentions involve commitment to conceptual action *types*, rather than particular bodily movements. This may seem slightly different from typical usage of the term “intention,” since most philosophers suppose that it is possible for me to form an intention to raise my right arm—a specific bodily movement. But I am trying to draw out a distinction between lower-level *immediate* action commands that directly result in and guide bodily movement—a particular raising of my arm, say—and action commands that are a degree removed from particular bodily movements. So as I use the term, the motor command that directly results in my raising my right arm this very moment is not an intention; it has no conceptual content.

For example, the wolf in our previous example might form an intention to hunt caribou. That intention will cause and guide the specific perceptual-level action commands that the wolf forms. Some of the lower-level action commands that the intention causes will cause and guide other more basic action commands, including specific motor commands. The wolf carrying out an intention to hunt caribou will not, assuming it makes no mistakes, form a perceptual action command to chase a rabbit, even though the rabbit constitutes food for the wolf and chasing it constitutes hunting. The wolf’s intention involves pure attribution—attributives like hunt and caribou that refer to kinds rather than particulars. The spider, by contrast, since it lacks a capacity for conceptual representation, only applies attributives that function to pick out particulars. This is true for its goals, as when it applies the attributive food to some particular object in its environment, and also for its actional representations, which always function to guide *particular*

context-bound bodily movements. So when the spider forms an action command to hunt, that representation instigates a particular series of bodily movements immediately.

It is important to note that the way in which intentions create distance between the representation and the action is *not* merely temporal. Even perception guides behavior diachronically, as the organism adjusts its movements in response to changes in its environment (or in its relation to its environment) throughout the course of some action. For instance, the spider might form an action command to jump on its prey that involves readying itself and waiting for the prey to come close, and then when the prey does so, jumping to the specific location that the prey actually occupies (or a prediction of where the prey will be at the termination of the jump, given the prey's current location and movement pattern). So the spider's action commands to hunt involve a range of possible movements and behaviors. What types all of these representations as representations of hunting are the goal (the prey) and the success conditions for the behavior (capturing and eating the prey). And at the level of perception there are differing degrees of generality, where many different specific bodily movements all constitute hunting behavior. The fact that an action command is a representation of hunting determines which specific bodily movements are mapped. The same thing happens with intentions, but the separation of the representation from the particular behaviors is of a different kind. The spider's action commands to hunt map out particular bodily movements as instances of hunting. The wolf's intention to hunt refers to hunting as an action type without necessarily applying the conceptual attributive to any particular instance of hunting behavior. The key difference between perceptual action commands and intentions is the same difference between perception and conception generally. Perceptual attributives are always and only applied to context-bound particulars; with perceptual-level action commands, actional attributives (like to

hunt) are applied to particular mappings of bodily movement. Conceptual attributives, on the other hand, can occur in pure attributions, where the property picked out by the attributive is itself the object of thought; with intentions, the action types themselves are the objects picked out by the actional attributives. This is what I mean by degrees of removal from particulars. This difference does have temporal implications, and an ability to think about action types themselves enables different opportunities for learning and planning. But the central distinction lies in the *kind* of representation.

It is also possible to form intentions that govern other intentions in a multi-layered structure. For instance, the wolf's intention to hunt might itself be governed by a broader intention to care for its pups. In that case, the intention to hunt is caused by and partly constitutes the intention to care for the pups. This layering of intentions allows for general intentions that require specification by more particular intentions. For example, I may intend to have a career, but that requires forming a more particular intention to take up this or that specific occupation. This point is important for understanding the hierarchical structure of human practical thought, which I consider in more detail in the next chapter. For now, note that a conceptual intention (for example, the wolf's intention to hunt) can be guided by a further intention (such as the wolf's intention to care for its pups).

Critical reason introduces yet another layer of practical thought. Critical reason is the capacity to represent intentions, beliefs, and reasons as such. For example, it is one thing to suppose that the wolf in our example can form an intention to hunt, and even perhaps form an intention to hunt that is guided by a broader intention to care for its young. But none of that entails that it can think about what its intentions are or form intentions that govern its other intentions *as such*. The wolf's intention to hunt might partly constitute a broader intention to care

for its young, and so it can have intentions that govern its other intentions in that way; but this need not imply that the wolf can think about what its intentions are and form different intentions in light of those reflections. Human beings can do this. We may be the only animals on earth with this capacity.¹²⁵

In human beings, a capacity for critical reason allows us to pursue a special kind of activity: *crafting a life*. This activity is set apart from other activities by the way that all other activities help constitute the activity of living a life. Thus, at least in the ideal case, all human intentions are formed under the broader intention to live a life, in the manner of multi-layered intentions described above. Now, in some sense, all organisms are engaged in the activity of living a life. And every thinking creature with a capacity for intention forms intentions to engage in activities (like hunting, mating, etc.) that constitute the living of a life suitable for the kind of organism that it is. But at least two human capacities set us apart from more primitive thinkers. First, a human being can form conceptual representations of her life as such. Thus, she can form an intention to live her life in a particular way, or under a particular conception, that thereby functions to guide all her other intentions, since the activities represented by those other intentions constitute her life. Second, a human being can represent and evaluate her concepts and intentions as such. That means that she can think about *how she conceives of her activities and how well she is pursuing them*. Moreover, she can understand her activities as being constituted by other activities, and so can reflect on how her intentions are related to each other.

I think that these two capacities—the capacity to represent the activity of living a life, and the capacity to represent activities as being constituted by other activities—are constitutive of critical practical reason. But there is another important element: Actions and activities are

¹²⁵ I am not committed very strongly to this hypothesis. I think it is at least plausible that other creatures—especially social mammals like apes, dolphins, and elephants—have critically rational capacities to some degree.

essentially teleological. They have success conditions. I think that a representation of an action or activity must be sensitive to those success conditions in some way in order to *be* a representation of an action or activity. For example, a representation of the action to eat must function to pick out *successful* eating, at least minimally, in order to count as a representation with that content. This point falls out of the nature of action commands and the evaluative representations that function to serve as goals within action commands. A basic action command represents a pattern of bodily movement, and functions to generate behavior that follows that pattern of movement. Like with all representations, the content of an action command is defined by its correctness conditions, and the correctness conditions of a basic action command just are the successful execution of the bodily movements represented. The same is true of conceptual representations of actions and activities. An intention to perform an action is a representation of that action in virtue of representing the (minimally) successful execution of the action. Similarly, a conceptual representation of an activity must contain some conceptualization of successful engagement in the activity. Without that content presenting successful execution or engagement, the representation would not be a representation of the action or activity in question.

Moreover, the patterns of activity represented in action commands are partly constituted by the evaluative representations that they take as goals. Some conceptual representations function constitutively to serve as goals for action commands. And conceptual representations can be developed over time so that they, in successful cases, are applied fluently in a more ample array of cases. Basic understanding, as I have explained it, just is fluency in applying a conceptual attributive veridically across an ample array of cases. So when an agent has a standing intention to engage in an activity, her basic understanding of that activity will determine what she sees as opportunity for engagement in that activity. And if she possesses a

conceptualization of the activity that is apt to serve as the goal of an intention, then she sees the activity as worth pursuing.

At this point we can dispel some of the supposed mysteriousness surrounding rational motivation. Since all action commands involve representations that function to serve as goals, it should come as no surprise that conceptual representations can move us to act just as lower-level representations can. The representations that function to be goals are evaluations. Just as the (proto-)evaluative representation as of food at the perceptual level functions to be the goal in perceptual-level conative states for eating behavior, conceptual evaluative representations can function to be goals in conceptual-level conative states (intentions). Rational conative states have the same general structure as lower-level conative states. At all levels, action commands cause and guide behavior and evaluations function to be goals in action commands. Thus, there is really no mystery as to how reason can motivate. Reason motivates by attuning us to actual values in the world in exactly the same way that perceptual representations motivate by attuning the individual to goods in its environment. Some values are simply goods relative to our own well-being, such as food and various pleasures. Other values are not relative at all, such as the value of individual persons. But in any case, motivation can be explained in terms of the conative function of evaluative (or proto-evaluative) representations.

Rational motivation has seemed mysterious to some philosophers because of a supposed dilemma faced by mainstream theories of practical reason. Sabine Döring (2007) calls this the “internalist dilemma”:

On an internalist account, it is part of the concept of a normative reason for action that it can also be the motive for which an action is done, and that is: that it can also explain an action. What makes a normative reason practical then is precisely that it can be both the reason why an action is right and the reason why it is done. To

accept internalism means to expect a psychological theory of motivation and explanation of action from the theory of practical reason: until it has been shown that we can be motivated by normative reasons, their validity has not been established. (364)¹²⁶

According to Döring (2007: 367), theories of practical reason fall into two main families, which she calls Kantian and Humean views. According to Kantian views, normative reasons must be mental states with representational content such that actions can be derived from rational, understood as inferential, processes. On a Kantian view, so-called “pure reason” alone can both motivate and explain (that is, justify) human actions. Humean views, on the other hand, insist that motivating mental states—usually called “desires”—are entirely arational, so that practical reasoning must always start from, and never result in, actual desires of the agent.¹²⁷ On a Humean view, reason alone never supplies motivation, but merely operates on motivational states that arise independently from it. Döring thinks that both of these views give rise to the dilemma:

With both the Humean and the Kantian view of practical reason normative and motivating reasons come to be seen as mutually exclusive. While the Humean cannot account for the normativity of practical reasons but commits himself to the incoherent claim that arational desires are capable of rationalizing actions, the Kantian fails for the opposite reason. Though the latter rightly points out that only states with a certain kind of content can enter into practical reasoning, he clings to the psychologically dubious postulate that pure reason has motivational force. (369)

Döring takes this dilemma to be an indictment of the belief-desire model of practical reason, which she thinks both the Kantian and Humean views accept.

¹²⁶ The term “internalism” has a wide variety of inconsistent uses in mainstream philosophy, which Döring acknowledges while nonetheless continuing to use the term in a restricted sense (2007: 367 fn. 5). I shall, however, avoid the term in my analysis because I think it causes needless confusion.

¹²⁷ A classic expression of this Humean position is in Williams (1979).

I linger on Döring's account because, for one thing, the view she settles on is certainly a version of the evaluative outlook view. But also, the way she sets up the dialectic illustrates a general problem that my view aims to avoid. In short, she proceeds with a distorted understanding of what motivation is, which causes her to misconstrue where the burden of argument lies. Essentially, she ends up granting too much to the Humean by conceiving of motivation as something special and *sui generis* that must be accounted for independently from the rest of philosophical psychology. In fairness, it may be that Kantian theorists make the same mistake, so Döring is responding to the dialectic as it stands in the literature.¹²⁸ But nonetheless it *is* a mistake, and it prevents a fuller understanding of representational agency. If Döring is right about the state of the dialectic (and I think she is), then the fact that this mistake is widespread makes addressing it all the more important.

I agree with Döring that the belief-desire model of practical reason, which insists (roughly) that rational actions are explained by pairing a *desire* that some state of affairs be the case with a *belief* that the action will bring about the desired state of affairs, is problematic.¹²⁹ I also agree that a complete rational psychology should make room for a wider variety of mental states, including phenomenal states such as emotions. I disagree with her conclusion that emotions are a necessary component of rationalizing explanations of action. But I have addressed these issues in previous chapters, so I will not rehash them here. Instead, I argue that a general account of representationally mediated action operates across the different kinds (or levels) of representational capacities, and this general account allows us fruitfully to reframe questions about motivation.

¹²⁸ I take no stand here on exegetical questions about interpreting Kant and Hume. I take Döring's labels to be reasonable enough given the contemporary dialectic, regardless of what the historical Kant and Hume actually thought.

¹²⁹ The classic expression of the belief-desire model of action comes from Davidson (1980).

To see how motivation is misconceived and how we can correct the misconception, it will be helpful to highlight two mistakes in how Döring sets up her argument. First, she simply grants without explanation the Humean claim that “an agent’s judgments cannot by themselves motivate any action, so [...] the Kantian theory of practical reason is psychologically inappropriate” (2007: 369). As best I can tell, she allows this concession to the Humean because she conflates two distinct claims made by the Kantian: (1) reason alone can motivate, and (2) reason should be what motivates in the ideal case. The first claim states simply that it is possible for thoughts—for instance, judgments about what ought to be done—to motivate action in the absence of desire, emotion, or any other more primitive or more phenomenally vivid mental state. The claim that such “purely rational” motivation is possible is warranted by reflection on our understanding of action and reason. Thus, this first claim enjoys a dialectical advantage: the burden of argument is clearly on whoever denies it. The second claim, however, is more ambitious, and possibly invites the critique that the Kantian position is “psychologically inappropriate.” This second claim insists that thoughts motivate in a way that is especially fundamental to rationally correct action, such that purely rational motivations are *superior* to actions motivated otherwise. But of course many human actions are primarily motivated by subrational capacities like emotion or desire, and yet such actions nonetheless often have rational explanations and can be justified. So this second claim has a stronger burden of argument.

Because Döring fails to make the distinction between the two Kantian claims, the psychological doubtfulness of the Kantian idea that purely rational motivation is superior to other kinds of motivation gets grafted onto the more modest claim that purely rational motivation is merely *possible*. But we need not accept both claims. I think that purely rational motivation—understood as action guided and caused by thoughts or judgments, and without the

aid of subrational states like desires or emotions—is indeed possible and happens commonly among humans.¹³⁰ People do things they do not feel like doing all the time. For most of us, it is easy to think of a time when we did some chore because we had good reason to do it, even though we did not want to do it. I remain agnostic, however, on the claim that purely rational motivation occupies some special place in the theory of practical reason, and I lean negative. More revealing are the cases in which rational capacities work in coordination with subrational capacities to cause and guide action. This kind of complex rationally mediated action is the paradigm of practical thought, in my view.

The second mistake in how Döring sets up her argument is that she tacitly imports the Humean idea that evaluation, and by extension motivation, is fundamentally different from non-evaluative representation. For example, when she discusses the emotion of fearing a snake, she writes: “I take it to be constitutive of your emotion that you feel fear of the snake, and that your fear represents the snake as being a certain way to you, say, as being dangerous. Because this is not a neutral representation of how things are, but an evaluation in light of your concerns, such as your concern to escape unharmed, your fear also has motivational force.” (2007: 372) Notice the contrast between “neutral representation of how things are” and “evaluation in light of your concerns.” But assuming that the fear is veridical, the fact that the snake is dangerous is just as much an element of “how things are” as anything else. Of course, the property dangerousness is relative to the nature of the individual whose fear it is, but that relativity does not make it any less an objective part of the world. Notice also that “concerns” seem to play a role suspiciously close to Humean arational desires. Internalizing this dichotomy between evaluative and

¹³⁰ Of course, some subrational capacities, such as perception, must be involved in carrying out any bodily action. The issue here is about what motivates and ultimately explains an action. No one thinks that intentions are formed and carried out without any input from perception and other subrational capacities.

non-evaluative states makes it seem strange and mysterious that a mere representation of the world could be motivating. So some other thing gets brought in to explain how representations could move the individual to act. But of course, now we must explain what a concern is and where it comes from. Thus, the theory ends up postponing the promised explanation.

My account avoids these problems by integrating the motivational force of representational states into their psychological function. Conative representations attribute properties by applying attributives that function to be goals in action commands, which are made veridical by the actions of the individual, actions which are caused and guided by those actional representations. This general form occurs across the various levels of representational sophistication, from subrational capacities like perception through critically rational intentions. Such an account lets us leave belief-desire psychology behind, while still explaining what that model was proposed to explain.

V. The Outward-Looking Nature of Practical Reason

My account of rational motivation can also explain how human beings are often moved to act without having a determinate sense of what they are trying to achieve. Consider what Brewer (2009) says about how people are moved to engage in dialectical activities. He claims that “one throws oneself into dialectical activities on the strength of an as-yet-indistinct intimation of their intrinsic value, hence without the benefit of a clear and definite representation of the mode of engagement that would properly consummate one’s desire” (37). So to put this in terms of my account of intentions, one can form an intention to pursue an activity, without having a definite understanding of exactly how that intention will shape one’s other intentions and actions. Rather,

adopting an intention to pursue an activity often involves a conception of the activity that evolves over time as one engages in the activity. It need not involve a representation of a definite state of affairs to be produced.

Now, Brewer presents this point in large part to combat the family of views that he calls “propositionalism,” which hold that agency, in one way or another, consists in producing desired states of affairs (2009: 17–24). Like Döring, Brewer rightly rejects belief-desire models of practical thought. But the idea that human agency often involves inchoate representations guiding our activities has a more profound import, one which Brewer presents in various ways, and which I hope to shed light on here from the perspective of an inquiry into representational mind. The point can be put as follows: *Human reason strives to achieve an ideal that cannot be articulated completely in terms of standards internal to human reason itself*. This claim requires unpacking.

Human rationality is ideally guided by an overarching intention: the intention to *live well*. And the content of this intention is not given by standards internal to rationality itself. Living well is the activity of trying to construct for oneself a *good life*. And in the fashion of all dialectical activities, the shape of this activity evolves over time as one engages in it. Engagement in the activity of living one’s life (at least when things go well) deepens one’s understanding of what constitutes living well and what constitutes a good life, and therefore what constitutes proper engagement in the activity of living well. And since the concepts and conceptualizations that make possible the overarching intention to live well and the overarching goal of having a good life are representations, they are subject to standards of correctness. And as with most representations, those standards are partly constituted by facts about the world outside the system in which the representations occur. Thus, in order to make sense of human

activities, we must have some grasp of what in fact has value—or put another way, we must employ a “substantive rather than merely formal conception of practical rationality” (Brewer 2009: 27).

Now, I think it is fair to say that not all human beings explicitly adopt an intention to live well. Perhaps most do not. On the other hand, there are many mundane kinds of practical deliberation that come very close to being attempts at formulating this kind of overarching intention. Consider how we commonly ask one another (and ourselves) “What are you going to do with your life?” Or “What kind of person do you want to be?” These sorts of questions, along with the kind of thinking involved in choosing a career or a spouse or a place to live, exhibit the kind of scope characteristic of an overarching intention, even if they do not necessarily involve explicit consideration of the question of what constitutes living well or having a good life. So even though the vast majority of human reasoning is mundane and narrow in its scope, we should not discount the rational sophistication of even relatively unreflective persons.

But if human practical reason answers to a standard external to rationality itself, what is that standard? It is to this question that I now turn. In a nutshell, the answer I offer is that human practical reason functions to respond correctly to what actually has worth, to what actually matters in the world.

Chapter 5

Meaning and Choiceworthiness: On the Shape of Practical Thought

According to Brewer (2009), practical thought is structured around the activity of living a good life. And for humans, the activity of living a life is a dialectical activity: a temporally extended activity carried forward by an ongoing understanding of the activity's value or point. On Brewer's account, the value of anything we do must be understood, at least in part, in terms of the place the activity has in a good human life. This formulation seems to invite the question of life's own value or point, or as Christine Korsgaard (1986) puts it, the question of what makes human life itself "worthy of choice" (492, 495). Now, perhaps it is unfair to demand an answer to such a large and difficult question in the course of theorizing about practical reason. Nonetheless, reflection on this question of the choiceworthiness of human lives sheds light on a great many other problems, and ethical theorists must wrestle with it. In what follows, I attempt to make progress on the problem, while not offering any definitive solutions to it, by exploring *meaningfulness* as a category of value within an evaluative outlook model of practical thought.

The notion of life's meaningfulness is distinctive as a kind of good in that it suggests a dimension of *independent value* that comes from outside the perspective of the person whose life it is.¹³¹ Moreover, the meaningfulness of a person's life is a matter of what kinds of activity the person is engaged in. Thus, several questions present themselves. First, what constitutes meaningfulness in life, and why does it seem to be an independent kind of value? Second, how is meaningfulness distinct from other kinds of value, such as prudence, morality, and happiness?

¹³¹ See Wolf (2012) for some arguments to this effect, especially pp. 24–33.

And finally, what makes some activities fonts of meaning, and what is the nature of the value that these activities have?

In this final chapter, I reflect on these questions with an eye toward the larger questions about the nature of practical reason and the choiceworthiness of human life itself. I emphasize the hierarchical structure of practical reason, which ought to be organized around the overarching end of meaningfulness. Meaningfulness, I suggest, is the value that life has when it is lived in proper relation to what actually matters. Thus, I aim to make good on my earlier claim that practical reason constitutively functions to respond to what matters. I also argue that the general value of meaningfulness in life cannot be understood apart from the value of particular meaningful projects, activities, and relationships. In other words, one cannot understand meaningfulness without a substantive understanding of what matters—of what has independent worth. This conclusion has implications for how to understand human well-being and the sources of value in human life.

In the first section, I summarize Susan Wolf's (2012) influential "fitting fulfillment" account of meaning in life, and I try to unify it with Brewer's (2009) view that dialectical activities are essential to living well. I argue that dialectical activity fleshes out some underdeveloped aspects of Wolf's account. In the second section, I argue that an evaluative outlook account of practical reason encourages an understanding of meaningfulness as the central organizing principle of practical thought. I also criticize Wolf's fitting fulfillment view, and offer a revised view that attends more carefully to the distinction between the representations that direct us toward meaningfulness and the value of meaningfulness itself. The third section discusses how practical thought is organized around general ends like living a meaningful life. I elaborate on the hierarchical structure of practical reasoning, and conclude that focusing on this

structure, although helpful, leaves some important questions about value unanswered. In particular, this focus makes it hard to see why valuing meaningful objects and activities only for the sake of meaningfulness is insufficient. Then I argue that the most fundamental value is not itself a general end of practical reason. Rather, the most important general end of practical reason—meaningfulness—is to live in proper relation to what is most fundamentally valuable. I conclude with some brief reflections on the overall choice-worthiness of human life, in light of these arguments.

I. Meaning and Dialectical Activity

Susan Wolf (2012) has presented an influential account of meaning in life. Wolf argues that meaningfulness is a category of value distinct from both morality and happiness (or “self-interest”), at least as those goods are normally understood in mainstream philosophy (1–7).¹³² As she presents it, meaningfulness is more central to typical human concerns than other goods because meaningfulness closely tracks the values that make human life worth living (2). Wolf explains meaningfulness as a product of “loving objects that are worthy of love and engaging with them in a positive way” (8), or as she puts the point elsewhere, “actively engaging in projects of worth” (26). Less vaguely, she argues that meaningfulness arises when *objectively* valuable activities or objects engender *subjective* feelings or experiences of fulfillment in the subject who engages in those activities or interacts with those objects (19–23). Importantly, Wolf arrives at this conclusion partly through observing that subjective feelings of fulfillment have a

¹³² I think the term “self-interest” is somewhat misleading here. Since whether or not moral uprightness makes one’s life better is itself an interesting philosophical question (and one discussed since the origins of the western philosophical tradition), taking morality and self-interest to be opposed categories of value without argument strikes me as question-begging, at least outside of peculiar contexts. Thus, I prefer to consider *happiness* as the value that is potentially opposed to morality, since happiness and morality are more clearly independent of each other.

“cognitive component that requires seeing the source or object of fulfillment as being, in some independent way, good or worthwhile” (24). In other words, if I may reinterpret Wolf’s view in the terms I have developed in this dissertation, meaningfulness in life requires that one form veridical evaluative representations with subjective, phenomenal modes of presentation. Put this way, it seems obvious that meaningfulness, whatever it is, has a subject-independent element, since representations are paradigmatically of objects and subject matters whose natures are what they are independently of the mental states that represent them. In the next section, I argue that Wolf’s account of meaningfulness needs to be revised to make room for the possibility that a life is meaningful but not fulfilling. These revisions rest on attending more carefully to two distinctions: the distinction between a representation’s mode of presentation and its veridicality conditions, and the distinction between the activity of living well and the end of living a good life. But before I get there, Wolf’s account needs further elaboration. She claims meaningfulness in life arises through “[loving] *engagement in projects of worth*” (26, my emphasis). But she does not say much about the nature of this engagement. It must be “positive” and “active” (6, 26), but a more specific development of these claims is needed. Fortunately, I think Brewer’s (2009) notion of dialectical activity helps fill this gap.

Dialectical activities are sustained by an evolving understanding of the value of the activity. This understanding is developed over time by continual engagement in the activity, and each iteration of the activity can shape new opportunities for the agent to build on in the future. Moreover, Brewer insists that activities with this dialectical structure play a fundamental role in human life, and although he does not discuss meaningfulness specifically, he thinks that dialectical activities are essential components of a *good* human life.¹³³ The good life as Brewer

¹³³ See Brewer (2009: 37–67) for a fuller discussion of dialectical activities and their place in a good human life. See also Chapter 1 of this dissertation.

understands it is almost certainly meaningful in Wolf's sense. Indeed, all of Wolf's examples of the kinds of projects apt to create meaning in life fit the mold of dialectical activity. And the notion of dialectical activity helps flesh out Wolf's notion of "positive engagement in projects of worth." Positive engagement is that which develops, or affords opportunities to develop, the agent's understanding of the value of the activity, and thus draws the agent into further engagement and into a closer, more fitting relationship with independently valuable things. And Wolf's emphasis on *love* as the fitting attitude toward objects, projects, and persons in an optimally meaningful life is paralleled in Brewer's writing, too, as Brewer frequently presents loving relationships with other persons—friendship, romantic intimacy, parenthood—as paradigmatic examples of dialectical activities.

Strictly speaking, the good life and the meaningful life are not identical. Of course, good lives are characteristically meaningful, and meaningful lives are better by virtue of their meaningfulness. But the goodness of a life constitutively incorporates the overall well-being and flourishing of the individual, whereas the meaningfulness of a life does not, although it indirectly contributes to the individual's well-being and partly constitutes the goodness of the individual's life. In my view, meaningfulness is the kind of goodness that a life has when it is lived in proper relation to that which matters in and of itself—that is, that which has independent worth. This view is reinforced by the fittingness of loving relationships as paradigms of meaningful activity, since other persons obviously have independent worth. If Brewer is correct in his characterization of dialectical activity, then dialectical activity is an essential part of responding correctly to what matters, and thus living in proper relation to what matters. Dialectical activity, then, can be understood as the general form of meaningful living.

One difference between Brewer's view and Wolf's lies in their characterization of the kind of value that dialectical and meaningful activities have. Brewer uses the term "inherent value" to describe the value that is apprehended over the course of one's engagement in dialectical activity (2009: 39–42). Wolf, on the other hand, focuses on "independent" and "objective" value (2012: 35, 41–48). Thus, in bringing these two views together, I must reconcile their disparate conceptions of value. I think they both have the same kind of value in mind, but it will take some work to bring that out. My first task, then, is to articulate an evaluative framework and vocabulary that captures the values in question and makes clearer some important distinctions.¹³⁴ An important upshot of this endeavor will be to take steps toward answering the question of the connection between valuable *activities* and valuable *objects and persons* on whom those activities might center.

I begin with the notion of inherent value that Brewer invokes. Brewer seems primarily concerned to distinguish the value of dialectical activities from mere instrumental value (2009: 39). So inherent value is *non-instrumental* value. This contrast is about the *source* of the value of an object or activity. Inherently valuable objects and activities are those whose value comes from their own nature. Instrumentally valuable objects and activities have a value that comes from constitutively independent, but causally connected, objects or outcomes. This distinction may be grasped by reflecting on the appropriate evaluative attitudes towards such things. It is appropriate to value an inherently valuable good for its own sake. Merely instrumental goods generally ought to be valued for the sake of something else.

Christine Korsgaard (1983) has argued that there are two important distinctions in this vicinity: (1) that between intrinsic value and extrinsic value, and (2) that between instrumental

¹³⁴ See also Chapter 1 of this dissertation, especially *Section III: On the Nature of Value*.

value and final value (169–173). Intrinsic goods are the source of their own value, while extrinsic goods get their value from some independent source. Final goods are valued for their own sake, while instrumental goods are valued for the sake of something else. These distinctions effectively separate the question of the source of a thing's value from the question of the appropriate evaluative attitude to take toward the thing.

Korsgaard is equivocal about whether the instrumental/final distinction is primarily concerned with how a thing is *in fact* valued or with how a thing *ought to be* valued. In fact, her argument depends on this equivocation, since she aims to show some things are extrinsically valuable precisely because they are in fact valued finally (for their own sake), and yet they are nonetheless correctly valued for their own sake. Her view—her version of constructivism—is that when a rational agent values something for its own sake, that valuing by the agent can be the source of the thing's actual value. Hence, the thing has extrinsic value, since the source of its value is external to it. But the fact that an agent values it for its own sake can at the same time make it appropriate or correct that the agent value the thing that way.¹³⁵ Insofar as Korsgaard's goal in making the two distinctions just described is to keep open the logical space for such an account, she succeeds. But by pushing the distinction between how something ought to be valued and how it is in fact valued into the background, she makes that view seem more plausible than it is.¹³⁶ I think the *prima facie* implausibility of the claim that valuing something makes it valuable

¹³⁵ Korsgaard argues for this view in various ways and in various places. See especially Korsgaard (1983, 1996, and 2003).

¹³⁶ Nonetheless, there is something right in the constructivist idea: rational choice does partly explain the role of various goods within the lives of particular persons. The rationally chosen projects and commitments of a particular person determine, in part, what that person has reason to do. But since value does not reduce to reason-giving properties, this does not mean that a person's rationally chosen projects and commitments determine what is good for her. See Garthoff (2010) for a discussion of how rational commitments contribute to a person's well-being. See Brewer (2009: 151–191) for an argument that the good does not reduce to reasons for action.

puts a fairly severe burden of proof on the constructivist. Paying attention to the representational elements of evaluation puts even further pressure on this kind of view.

For now, I want to expand upon Korsgaard's distinction to help understand the kind of value that dialectical activities have. First, I prefer the labels "independent" and "dependent" to Korsgaard's "intrinsic" and "extrinsic," since the point of the distinction is whether or not a thing's value depends upon, or is grounded in, the value of something else. Thus, I call anything that gets its value from its own nature *independently* valuable. Anything that gets its value from something other than its own nature is *dependently* valuable.¹³⁷ Second, instrumental value should be understood as a kind of dependent value that has a *productive causal* connection to the object upon which the value depends. But dependent value can also be *constitutive*. Constitutive dependent goods depend for their value on the fact that they partially or potentially constitute something else valuable, usually through instantiation or composition.

The question of an agent's subjective attitude toward a value is separate. An agent can, correctly or incorrectly, value something for its own sake or for the sake of something else. In general, it is correct to value an independently valuable thing for its own sake, since the source of the value is the thing's own nature. An instrumental dependent good should be valued for the sake of the valuable outcome it causes.¹³⁸ A constitutive dependent good should be valued for its own sake when the thing it constitutes is independently valuable (and it constitutes that valuable thing in the right kind of way). Some goods, like meaningfulness, are general and so are

¹³⁷ My notion of independent value is similar to G. E. Moore's (1922) concept of "intrinsic value" in that we both identify the kind of value in question as being connected to a thing's *nature*. Unlike Moore, however, I have no particular theory of "intrinsic nature." I think that the notion of a nature or essence is pretheoretically available and needs no vindication by theory to be useful. See Burge (2010a: 57–59).

¹³⁸ Of course, a good can be independently valuable *and* dependently valuable at the same time, which complicates the evaluative attitude that one ought to have towards it. For example, a good meal can be good for its own sake, but also a means to good health.

constituted (and multiply realizable) by various objects, activities, and relations. When these general goods have independent value, one can only achieve them by valuing their particular constituting constituents, and those constituents must be valued for their own sake in order to successfully constitute the more general good. For example, it is probably true that one cannot find meaning in life without being in some kind of loving relationship with others (at least outside of highly unusual circumstances). But if one values one's loving relationships only for the sake of one's own meaningful life, those relationships lose their power to constitute meaningfulness (probably because they thereby cease to be genuine loving relationships).¹³⁹ I shall say more about this point shortly. For now, I return to dialectical activities.

With these distinctions in hand, we can begin to understand the kinds of value that dialectical activities have. At first blush, it seems obvious that dialectical activities have constitutive dependent value. That is, dialectical activities aim at various general goods (such as having a career, a friendship, or a meaningful project) that together constitute a good or meaningful life. Living a good life is the most general good that can serve as an organizing principle for all the other goods that one might pursue. Thus, the human good involves what Garthoff (2010: 709) calls a "nesting structure" wherein the most general aim a person can have, living a good life, is constituted by less general aims, like having a friendship, which in turn are constituted by less general aims, like having dinner with this particular friend, and so forth. The notion of a dialectical activity, then, is a specification of the kind of activity that constitutes the

¹³⁹ This is an obvious analog to the familiar "paradox of happiness." But I am making more than just a psychological point. Rather, I want to emphasize the constitutive element of the claim. The point is not just that aiming at something like pleasure can make it less likely that you will achieve it; more than that, the reason for the paradox, at least in certain cases, is that some goods *only exist* as instantiations of the more particular goods that constitute them. For example, if Brewer (2009: 108–150) is right that some pleasures *just are* experiences of "unreserved activity," then one cannot achieve those pleasures without engaging fully (unreservedly) in such an activity. This engagement requires attending to and aiming at the specific activity and its goals, not at the pleasure itself (which is just the phenomenal component of experiencing engagement in the activity). See also Garthoff (2010).

human good. Moreover, since the activity of living a good life is itself dialectical, the structure of dialectical activity helps explain the overall nested structure of the human good. Dialectical activities unify what might otherwise seem to be disparate and unconnected actions.¹⁴⁰ For example, the activity of friendship brings together a wide array of seemingly unrelated activities—helping my friend move, giving her a ride to the airport, preparing a meal together, going for a walk, playing a game, and so forth—into a single continuous pursuit.

This picture seems correct so far as it goes, but there are important questions that it does not satisfactorily address. For one thing, it leaves untouched the question of what makes human life itself choiceworthy. And this question looms large if our goal is to explain fully the value of dialectical activities. The explanation on offer grounds the value of any activity in how the activity contributes to a good human life, so a full account of that value ought to explain something of why good human lives matter. Moreover, the value of many activities seems to be grounded not just in their contribution to living a good life, but in the independently valuable objects or persons around which the activities revolve. Activities of friendship are a prime example. It seems implausible that the value of friendship in general—or the value of living a good life, for that matter—explains the value of particular friends. Thus, more needs to be said about the nature of independently valuable entities like persons, and about how their value contributes to the meaningfulness of our activities, projects, and relationships. It may be illuminating and necessary to explain how an activity, project, or relationship fits into the larger picture of a good or meaningful life. But that explanation will not tell the whole story, since it will not be able to explain what makes human life itself worth living. Nor will it be able to explain why a meaningful life is preferable to a meaningless but happy one, if indeed that is the

¹⁴⁰ See Brewer (2009: 37–67).

case. To begin to formulate answers to these questions, I must return to the evaluative outlook view of practical reason.

In the next section, I summarize what I have said in previous chapters about the evaluative outlook view, applying insights from that view to Wolf's account of meaningfulness. I present meaningfulness as a kind of excellence of human life. Specifically, a meaningful life is a life lived in proper relation to that which has independent worth.

II. Evaluative Outlooks and Fundamental Values

My evaluative outlook view of practical reason builds on Brewer's (2009) notion of dialectical activity. Brewer's explanation of dialectical activity depends upon his account of desires as "apprehensions of value" (24–34). According to his account, desires are both representational and essentially phenomenal. Desires represent their object as good in some way, but they do so with a phenomenal vividness that distinguishes them from other more removed evaluative representations, like evaluative judgments. Brewer proposes this notion of desire to explain how a dialectical activity draws agents in through their initially inchoate grasp of the activity's value, and also how one's engagement in a dialectical activity can lead to different and more refined ways of engagement as one's understanding of the activity's value deepens. Since desires are perception-like, they allow us access to values that we do not yet fully understand. Dialectical activities are valuable in a way that cannot be fully grasped short of engaging in the activities themselves, and yet we must be capable of grasping their value in some way such that we are drawn to the activity in the first place. Otherwise, our initial engagement in any specific activity, or our initial taking up of a specific project, would be arbitrary. Any sufficiently similar activity or project would do. But upon reflection, this is not how we experience our valuable

projects and activities. Rather, we feel drawn to specific people, objects, and activities by the granular particularities of those things. Moreover, we need an account that makes sense of the fact that we do not begin our lives as critical reasoners, but develop capacities for critical reason and understanding over time. Conceiving of desire as perception-like helps explain these things, since perception occurs through sub-personal processes that are not reflectively accessible, and perceptions are always caused by and responsive to particulars.

Yet desires must also be sensitive to conceptual understanding such that they can be shaped and redirected over time as the agent's understanding of an activity and its value develops through her engagement in the activity. Thus, another aspect of the perception-like nature of desire is the potential for desires to be "conceptually structured" (Brewer 2009: 24–31). When perceptual representations enter the space of conscious awareness, they are often already shaped by conceptualization. Thus, the sub-conscious processes that generate our perceptual representations include or are supplemented by conceptual representational processes that apply conceptual attributives on the basis of perceptual inputs.¹⁴¹ If desires, too, are shaped by our conceptualizations of their objects, and if those conceptualizations are refined and developed through engagement with the objects of the desires, then we should expect that our understanding of an activity and its internal goods may be deepened by experience and may shape our future desires for engagement in the activity.

It should be emphasized that on this account the value of an activity is an essential part of its nature (Brewer 2009: 38–40). Activities are essentially teleological, so understanding what the activity *is* involves at least some grasp of the activity's constitutive goals, purposes, and ideals. Seeing this point clearly helps bridge the gap between the epistemic premises of Brewer's

¹⁴¹ See Chapter 2 of this dissertation.

argument and his ontological conclusions. For Brewer, the value of a dialectical activity can only be understood alongside an understanding of the activity's place in a good human life (2009: 41). And this epistemic relationship between valuable activity and the good life goes both ways: one understands valuable activity through one's grasp of the good life, and one understands the good life through one's grasp of valuable activities (41). One cannot fully grasp the value of a dialectical activity outside of engaging in the activity because understanding the value of such an activity involves an ineliminable phenomenal component, which comes from the experience of desiring and engaging in the activity (45–49). The thrust of Brewer's reasoning, as I understand him, is that these epistemic facts are *best explained* by ontological conclusions about the nature of desire, dialectical activity, the good human life, and the objective reality of value. Recall the seven key claims of the evaluative outlook model that I presented in the first chapter:

- 1) Some mental states are evaluative representations with ineliminably evaluative representational content.
- 2) Any successful explanation of a rational action must appeal to evaluative representations.
- 3) Many evaluative representations are products of subrational processes and capacities.
- 4) At least in humans, evaluative representations function to be elements in an overall evaluative understanding, or conception of the good.
- 5) Evaluative representations by nature have motivational efficacy.
- 6) Some evaluative representations are partly constituted by their phenomenal quality.

- 7) The phenomenal quality of an evaluative representation (if it has phenomenal quality) partly explains its motivational efficacy.

The first five of these claims are part of any evaluative outlook account of practical reason. The last two come from Brewer's particular elaboration of the evaluative outlook view. For Brewer, human practical thought is structured by dialectical processes, and these processes involve grasping actual goods whose natures are independent of any one particular point of view. These dialectical processes are themselves structured by the overarching aim of living a good life. But this general aim is always constituted by particular relationships, projects, commitments, and activities. One cannot understand the nature of the good life without attending to these more particular goods. Importantly, the activities that are able to constitute the good life for any particular person are deeply contingent, and the value those activities have is grounded in particularities of the activities and their objects. In order for an activity to constitute the good life for an individual, the individual must be able to grasp the specific qualities that make the activity valuable. And the ideals and aims of an activity, and thus also its value, are part of the nature of that activity, so understanding what an activity is must include some grasp of the activity's value and point. Finally, the only way to grasp fully the specific qualities of an activity that make it valuable is through qualitative phenomenology. And this phenomenology is part of what it is to desire engagement in the activity, which is to see the activity as good in some way. The evaluative content of desire, at least in humans, has an ineliminable phenomenal aspect. I suggest that this aspect is part of the desire's mode of presentation.

Once again, I don't want to overstate the importance of the phenomenology of desire. One can understand and appreciate quite a lot about the value of an activity without desiring to engage in it oneself. And while the phenomenal component of evaluative understanding is

important, it is not a basic requirement of *all* evaluative understanding. Rather, the phenomenal mode of presentation characteristic of desire enables a richness of evaluative understanding, which explains one important source of motivation.

But considering phenomenology brings us back to Susan Wolf's account of meaning in life. In Wolf's view, meaningfulness is constituted by the active pursuit of worthwhile activities, but only when that pursuit engenders positive qualitative subjective feelings or experiences of fulfillment.¹⁴² I think the subjective experience or feeling of fulfillment is a phenomenal aspect of a correct representation of the actual worth of one's activities. I also think we should understand the kind of active engagement that Wolf's account requires as a species of the same dialectical process that Brewer articulates: desire, followed by experience and deepened understanding, followed by reshaped desire, and so on. This picture makes sense of both the objective and subjective components of meaningfulness required by Wolf's account. On Brewer's account, mature desires are informed by a correct evaluative understanding, but they are not reducible to mere correct representations. They also have a phenomenal component that is an ineliminable aspect of their evaluative content. Thus, we can understand fulfillment (the subjective aspect of living a meaningful life) as the phenomenal feeling that arises out of the satisfaction of vivid desires, and meaningfulness occurs when desires that engender feelings of fulfillment are for things that in fact have the value that the desires represent them as having. At least, that is how Wolf's account would have it. But I think this is not quite right. To see why, we must consider the possibility of a life that is meaningful but not fulfilling.

¹⁴² Wolf (2012: 109–115) clarifies that perhaps there is no one singular kind of positive subjective experience characteristic of living a meaningful life. Feelings of fulfillment may be one kind of such positive experience, but she concedes there may well be others. For ease of exposition, I set this complication aside and use the term “fulfillment” to indicate whatever range of positive subjective feelings or experiences are constitutive of meaningfulness in Wolf's view.

It seems correct upon reflection that a person's activities can in fact be meaningful even if she does not get any subjective sense of fulfillment from them. For example, consider a dogged researcher whose work makes significant contributions to scientific knowledge, but who pursues her research entirely for extrinsic reasons, perhaps to please her demanding parents. Suppose this person gets no joy, no sense of fulfillment, from her work. Given that she in fact makes important contributions to science, I am inclined to think that her life is meaningful nonetheless. Of course, Wolf is correct that there is *something* important missing from the life of such a person. But perhaps we should not think that this missing piece is constitutive of meaningfulness. After all, the question of how meaningfulness makes my life better *for me* seems to be a live one. Compare our hypothetical unfulfilled researcher with a second person who conducts research of equal scientific value, but who does so out of a genuine love for the work and who thereby feels fulfilled by her scientific career. One might think of these two characters that their lives are equally meaningful but unequally fulfilling. If such a characterization is not confused, then we ought to think fulfillment is not a necessary constitutive component of meaningfulness. In short, meaningfulness and fulfillment come apart.¹⁴³

In her defense of the idea that meaningfulness includes an objective component, Wolf invokes the idea that human rationality demands that a person's actions and judgments be justified from a perspective larger than just that of the individual (2012: 27–30, 53–59). But this line of thought makes it hard to understand why meaningfulness should have a subjective component. Why does it matter, from a depersonalized point of view, that I feel fulfilled as long as my life is in fact devoted to worthy projects and pursuits? Thus, one might think that there need be no subjective dimension to meaningfulness at all. I find this objection to Wolf's "fitting

¹⁴³ Even Wolf herself seems tacitly to accept the naturalness of this way of thinking when she writes of "meaningfulness *as well as*... happiness or fulfillment" (128, my emphasis).

fulfillment view” of meaningfulness compelling, and so I propose a slightly revised version of the view. I have argued that some evaluative representations constitutively have phenomenal modes of presentation, and that the phenomenal quality of these modes of presentation iconically mirrors actual qualitative and evaluative features of reality. I suggest that we understand fulfillment as the phenomenal mode of presentation of desires for meaningful objects and activities. Understood this way, Wolf’s account is really an account of *fulfillment* as an aspect of a certain kind of representational content, rather than an account of meaningfulness itself. In fact, by distinguishing more carefully between the representational content expressed through feelings of fulfillment and the value that such content represents, the view has even greater explanatory power.

Wolf does not seem to agree that this line of critique is the most compelling reason to reject her view. She spends much more energy defending her view against the exact opposite line of attack, which claims subjective fulfillment is the only requirement for meaning, and thus, there is no objective criterion (2012: 36–48, 62–63, 119–132). Perhaps she supposes that most of her readers are likely to accept the subjective criterion but reject the objective one. She may be right about that. Nonetheless, I think the connection between meaningfulness and fulfillment is not correctly developed by Wolf. Fulfillment is surely connected to meaningfulness, but it is not a necessary constitutive condition of it. In fact, by keeping our eye on the nature of representation we can see that the subjective view of meaningfulness, which says that any activity or way of life that makes someone feel fulfilled is thereby meaningful, is implausible on its face because it cannot make sense of the representational component of feelings of fulfillment. Wolf herself is sensitive to this concern (2012: 24). But she does not really leverage the point to its full effect against the subjectivist critic. The subjectivist must either deny that

feelings of fulfillment have representational content, or claim that feelings of fulfillment are a very strange sort of representation, which are somehow always veridical by virtue of simply existing. I think that reflection on experiences of fulfillment is enough to warrant rejection of the first claim. Feelings of fulfillment are not brute pleasures. We do not pick them out by some generic positive valence or subjective quality common to them all. Rather, we pick them out, at least in part, by what they are *about*. As for the second subjectivist claim, there are some representations that do, by virtue of their peculiar content, guarantee their own veridicality. For example, my occurrent belief I am thinking cannot be false simply because having any occurrent belief at all is sufficient to guarantee that I am thinking. But such guaranteed representations always have some sort of reflexive or tautological content that explains why their existence is sufficient for their correctness. The representational content of feelings of fulfillment is not like that. It is outward-looking and substantive.

So either way, the burden of proof is squarely on subjectivists about meaningfulness. They must provide a strong argument to get us to jettison our understanding of the representational nature of fulfillment. I see little promise that a sufficiently compelling argument will be produced, especially given the explanatory power of the evaluative outlook view as an objectivist alternative. This reasoning generalizes to subjectivist views across the board. I articulated a version of this argument against Korsgaardian constructivism above. The general point is that once we recognize that evaluative representations are characteristically about the world, just like non-evaluative representations, fully subjectivist accounts of value lose much of their appeal, since they must either deny that there are any evaluative representations or take on radically revisionist stances toward the nature of representation itself.

An evaluative outlook view does much better at explaining the connection between the meaningfulness (the value) and fulfillment (the feeling) than any kind of subjectivism. I argued above that reflecting on the possibility of equally meaningful but unequally fulfilling lives warrants the conclusion that meaningfulness and fulfillment come apart. Now, one way to explain this is to distinguish between the activity of living well and the end of having a good (or meaningful) life. Meaningfulness may be best understood as a value that applies to a life in its totality. Thus, meaningful activities are those that are apt to make one's life meaningful overall, and there is no necessary subjective criterion on the overall meaningfulness of one's life. On the other hand, living well is an ongoing activity, which may have a necessary subjective element. Fulfillment could be a constitutive component of living well. Plausibly, the person who does meaningful work but finds no fulfillment in it is living less well than the person who does equally meaningful work but is fulfilled by it. Brewer's account of desires as phenomenally vivid representations of value, which are motivating (in part) by virtue of their phenomenal quality and reflective of a thorough evaluative understanding, can explain this idea further.

If we suppose that phenomenal feelings of fulfillment are the modes of presentation of a certain kind of evaluative representation, the connection between fulfillment and meaning comes into view. If one devotes oneself to a worthwhile project, but fails to find fulfillment in that project, that may indicate a failure to grasp the project's value fully. A person in such a situation would have at best a removed, abstract sense of the value of her project. A fuller understanding of the project's worth would involve a more direct, and therefore more phenomenally vivid, apprehension of the project's value. Recall also that understanding involves a robust integration of conceptual representational abilities into one's representational psychology (see Chapter 4 of this dissertation). So an understanding of something's value that is not incorporated into one's

practical representations (like desires) is an incomplete understanding. Some activities, projects, and relationships are meaningful, and this kind of value might be correctly represented in the subjective mode characteristic of fulfillment. Thus, Wolf's arguments might be better characterized as an account of *fulfillment* and not meaningfulness. Fulfillment, as I have presented it, is a constitutively qualitative subjective attitude, which is also constitutively representational, and thus partly constituted by correctness conditions. Whether or not an activity is fulfilling depends upon that activity in fact having the value that feelings of fulfillment present it as having—namely, meaningfulness. I think reflection on the difference between reality and representation—on the difference between *being* fulfilled and *feeling* fulfilled, in this case—warrant this conclusion. This point explains the difference between lives that upon reflection we judge to be equally meaningful but unequally fulfilling. It also explains our resistance to calling morally repugnant activities fulfilling, even if those activities feel fulfilling to those engaged in them.

Now, some finer points of this view deserve mention. First, the ways in which meaningfulness and fulfillment come apart may be myriad and complex. Even if we grant that the goal of a project or activity is sufficient to make it meaningful, there are many ways for the goal of an activity, or the source of its value, to become divorced from the daily grind of engaging in it. My previous example of scientific research, for instance, may have been too quick to assume that such research is fulfilling, or even that such work is meaningful. Suppose that contributing to scientific knowledge is, in the abstract, a meaningful enterprise (a claim that is not beyond dispute). Even so, that is no guarantee that the day to day activities of even the most devoted and clear-eyed scientist will be sources of fulfillment. If we consider the ways that such a career, especially within the apparatus of modern institutions of research, might be largely

composed of tedium, political maneuvering, and perverse incentives, we may come to expect many if not most actual scientists to lose sight of what makes their activities worthwhile in the first place. And lest we suspect that the bloated complexity of modern bureaucracies are entirely to blame, consider another more intimate example. Imagine a person devoted to caring for a dependent loved one. It is easy to see such care as meaningful, but again the day to day frustrations of such a task are equally apparent, even if we suppose for the sake of the example that the individual has ample time and resources to devote to the task. And no doubt there are important differences between an activity that delivers diminishing feelings of fulfillment as time goes on and an activity that never delivers feelings of fulfillment, and between activities that ought to deliver fulfillment but do not and those that are unfulfilling because they in fact lack meaning. These considerations suggest that meaningfulness without fulfillment may be an unfortunately common occurrence. Moreover, the elusiveness of fulfillment may be evidence against the meaningfulness of many activities that we otherwise assume to be meaningful.¹⁴⁴

This brings me to a second caveat in my account of meaningfulness: the meaningfulness of an activity may be quite difficult to assess, even if the presence or absence of fulfillment is clear. Thus, I want to stress that my arguments do not imply any judgments about the worthiness of particular activities.¹⁴⁵ Just because I do not find certain pursuits, like lawn-mower racing, intelligible as fulfilling activities, I need not thereby conclude that they are not meaningful. (But if they are meaningful, the source of that value is not *merely* the fact that people find them enjoyable or fulfilling.) Of course, it should come as no surprise that those who do not engage in the activity might not be able to see its value clearly from the outside. It is only by desiring the

¹⁴⁴ I would like to thank Claire Darteze, Georgi Gardiner, Jon Garthoff, Linh Mac, and Troy Pollidori for bringing some of these issues to my attention.

¹⁴⁵ Wolf (2012: 39–48) emphasizes the same point. The example of lawn-mower racing is hers.

activity and actually engaging in it that one can fully understand and appreciate its value. Nonetheless, it is possible to see value in an activity that is not really there. Thus, it is often appropriate and always intelligible to question the value of one's own activities. This questioning does not take the form of asking myself whether I really find the activity fulfilling, but rather asking whether the activity really has the value that I see in it. If we take fulfillment to be a species of pleasure, this amounts to asking whether the pleasure and satisfaction that I find in the activity are appropriate. Of course, in itself pleasure does not have correctness conditions. Like pain, pleasure is constituted by its subjective presentation. But if pleasure is, at least sometimes, an aspect of the mode of presentation of an evaluative representation, then the appropriateness of pleasure can be assessed by assessing the correctness of that representation, of which the pleasure is an ineliminable aspect.¹⁴⁶ Thus, when I ask myself whether something I take pleasure in is really good, I am asking whether I *ought* to take pleasure in it, and at least sometimes, that “ought” should be understood as a norm of veridicality applied to the representation for which pleasure is the mode of presentation.¹⁴⁷

The availability of this kind of questioning of one's projects and commitments suggests the organizing role that general goods, like living a meaningful life, play in practical thought. On

¹⁴⁶ See Brewer (2009: 108–150) for an extended account of pleasure as “unreserved activity” rather than as a phenomenal feeling or experience. I do not think Brewer's account captures all forms of pleasure, but I do think it is plausible as an explanation of one particular kind of human pleasure. Understood that way, my suggestion that pleasure can be part of the mode of presentation of an evaluative representation may reinforce that view.

¹⁴⁷ I think there is an important difference between asking whether something I take pleasure in is really good and asking whether it is really good *that* I take pleasure in something. The former focuses on the source of the pleasure and highlights the epistemic and representational role of pleasure. The latter focuses on the wider context, highlighting the way that pleasure can be good or bad, all things considered, in a way that does not specifically depend on the source or nature of the pleasure itself. Consider, for example, two ways one might wonder about the pleasure of eating a decadent piece of cake. On the one hand, one might wonder whether one ought to take pleasure in eating the cake because one questions the cake's objective quality. On the other hand, one might be more concerned with how the (perhaps entirely appropriate) pleasure one takes in eating the cake might impact one's health in the long term, given other factors about oneself, such as weakness of will. The first case involves a concern about whether the cake is really as good as the pleasure suggests that it is. The second case does not question the cake's goodness, but rather the goodness of the life in which the pleasure of eating the cake is situated.

the one hand, it is true that general goods help explain the value of the goods that constitute them. For example, my friendship with a particular person is good (for me and my friend) in part because of the role friendship generally plays in a meaningful life. But on the other hand, the value of general goods is also partly *explained by* the particular goods that constitute them. Having a meaningful life is valuable in part because it is constituted by worthwhile things, like particular friendships. This point is warranted by reflection on the elusiveness of general goods within our practical thought. Usually, it is only by keeping our attention on the specific activities that constitute a meaningful life, for example, that we are able to realize that general good. Nonetheless, the way those specific activities fit into the larger project of living a meaningful life shapes and informs how we ought to pursue them. Recall Brewer's observation that there is a two-way epistemic relationship between dialectical activities and living well: we understand the value of particular activities in part by understanding their place in a good life, but we also understand the good life in part by understanding the value of particular worthwhile activities (2009: 41).¹⁴⁸ Thus, it appears that neither general goods nor the particular goods that constitute them have overall explanatory priority. Defending this claim requires further reflection on the overall structure of practical thought.

III. The Structure and Content of Practical Reason

So practical thought is organized around general goods, like living a meaningful life, that are constituted by less general goods. Human practical reason, then, has a hierarchical, or “nested,” structure (Richardson 1997, 2004; Garthoff 2010). I think that meaningfulness occupies a special place in the hierarchy of practical thought, despite not being the most

¹⁴⁸ It is worth noting that I use the term “epistemic” broadly to include all aspects of conceptual representational psychology, so not just knowledge, but belief, conceptualization, and understanding as well.

comprehensive good toward which practical reason might aim (that distinction goes to the good life). But we cannot understand meaningfulness as a good without attending to other values whose natures are not grounded in the structure of practical reason. In this section, I elaborate these ideas. First, I describe the general structure of practical thought as a hierarchy of ends in which the most comprehensive ends are constituted (and multiply realizable) by more particular ends. Then I reflect on the particularity of value, and I suggest that meaningfulness occupies a special place at the top of the hierarchy of practical thought because meaning in life just is the proper relation between the individual and what actually matters: the particular things in the world with independent worth.

Some ends depend for their value on the fact that they constitute other independently valuable goods, but they can only constitute those independently valuable goods if the agent treats them as valuable for their own sakes (Schmidtz 1994; Richardson 1997; Garthoff 2010). Thus there are independently valuable ends that stand in need of being constituted by more particular ends, which must themselves be valued for their own sakes in order to constitute the more general ends. These ends are general in the sense that they are multiply realizable and must be constituted by other more particular goods. For example, no one has a career except by having some particular career, such as being a lawyer, a teacher, a nurse. And some ends can both serve to constitute more general ends and themselves be constituted by more particular ends. For example, being a teacher is one way to constitute having a career, and having a career is one way to constitute having a meaningful life project.

We can leverage this hierarchical account of practical reason to explain why adopting something as an end can seem to imbue that end with value that it otherwise would lack (Schmidtz 1994, Garthoff 2010). The idea is that the adoption of an end as a constitutive means

for achieving an actually valuable final end “sustains” the value of that constitutive means and makes it correct for the agent to treat that end as valuable for its own sake (Garthoff 2010: 705–706). In this way, we can maintain the constructivist insight that our rational choices and reasonable desires shape what is in fact good for us (Korsgaard 1983, 1996, 2003). But we can at the same time avoid the problematic constructivist conclusions that there are no substantive constraints on what can be valuable and that rational desires can create value out of nothing. We can maintain that desiring something in no way makes that thing valuable if it does not already have value independently of the agent’s desire. Yet we can also explain how the correctness of an agent’s *treating* something as valuable depends in part on that agent’s actually adopted aims and projects.

An end is anything that an agent treats as valuable, and (for humans) treating something as valuable is to accept it, either tacitly or explicitly, as a “volitional constraint” (Richardson 1997; Garthoff 2010). Recall that conative states contain *goals*, which are representational attributions that function to play a targeting role in *commands*, which are representations that function to cause and guide specific actions (see Chapter 2 of this dissertation, especially *Section II: Evaluative Representational Content*). Recall also that conceptual action commands are *intentions*, and action commands, especially intentions, often operate at different levels of generality (see Chapter 4 of this dissertation, especially *Section IV: Intention and Rational Motivation*). Thus, a more comprehensive intention can be composed of or instantiated by less comprehensive intentions. To draw on the previous example, one can intend to have a career, a general intention which is instantiated in the more particular intention to become a teacher, which is composed of various even more particular intentions, such as enrolling in a teacher training program, and so forth. Similarly, one can intend to form and maintain friendships, which

is instantiated in the more particular intention to befriend a particular person, which includes various intentions to spend time with that person, and so on. I shall reserve the term *end* to refer to the goals of actually adopted intentions.

Now, goals function to pick out valuable things—activities, objects, and kinds with valuable properties. In the most primitive cases, the values picked out as goals are goods relative to the functions, flourishing, and well-being of the individual. For example, a perceptual-level representation of food picks out a goal for commands to eat, and other related activities. And the value of that goal is a good relative to the biological functions (nutrition) and biological flourishing of the individual. In a similar way, the goal of having a career is a good relative to the social and psychological flourishing of an individual person in a society like ours. Thus, when one adopts the end of having a career, one might do so in part *for the sake of* one's own flourishing. This “for the sake of” relation only operates explicitly at the level of thought because it comes out of the individual's conception of the goal towards which the intention aims.

Intentions function to fit within a network of inferentially connected ends and other intentions. Thus, if one intends to have a career, and yet also intends to eschew having a career to care full time for a sick family member, one must reconcile these two conflicting intentions. Intentions are also sensitive to the individual's overall understanding (or conception) of the value of the end. If one intends to have a career not for the sake of one's own flourishing but only for the sake of satisfying the expectations of one's parents, then that intention will fit with and shape one's other intentions differently.¹⁴⁹

Critically rational intentions have a further function to be sensitive to values beyond mere goods relative to function, flourishing, and well-being. This is because critically rational

¹⁴⁹ I think this latter point holds even for the intentions of individuals without critical reasoning capacities. See the example of the dog in the next paragraph.

intentions must be fit to survive reflective evaluation *as* rational intentions. With critical reason comes the capacity to conceive of values as such. Suppose, for example, a dog has a capacity for conceptual thought but no metacognitive or critical reasoning capacities. The dog might form an intention not to pee inside the house for the sake of avoiding his owner's anger. (I am assuming in this example that the dog's conception of his owner's anger does not include any representations of psychological states as such.) That intention would be functionally different from an intention not to pee inside the house for the sake of keeping his living space unsoiled. Of course, he could form the intention for both reasons. But the point is that the dog's conception of the value of his ends will shape his conative psychology even in the absence of an ability to reason about and evaluate his reasons as such. Indeed, it is how the end figures in the dog's overall psychology that determines that the end is adopted for the sake of this or that. The dog need have no concept of value, goodness, or well-functioning in order for him to adopt an end for the sake of his owner's happiness. It is enough that avoiding his owner's displeasure in fact operates as a reason guiding his behavior in the right kind of way (which is not to say that it must *always* guide his behavior—one can act against one's adopted ends and still intend to pursue those ends.) Contrast this case with that of a human being with a mature capacity for critical reasoning. The human being can adopt an intention to pursue a career for the sake of her own flourishing, which means that she takes her own flourishing to matter. Upon reflection, then, she must either endorse her own flourishing as a good reason for adopting the end or revise the end, either by abandoning it or finding another sufficient reason to pursue it. Thus, in order for her end of having a career to be justified, it must be the case both that having a career contributes to her flourishing (since it is for the sake of her flourishing that she has adopted the end), *and* that her flourishing matters independently (since her flourishing must be a sufficient reason to adopt

the end). Of course, her flourishing does indeed matter, since all persons have independent worth. But the point is that critically rational ends function to be sensitive to all sorts of value, while perceptual and merely rational goals need not be sensitive to values outside of goods relative to individual functioning, flourishing, and well-being.

Since treating something as valuable, adopting it as a volitional constraint, is subject to correctness conditions, reflection on the nature of ends informs reflection on the nature of value and vice versa. To suppose it is correct or appropriate for an agent to value something for its own sake invites reflection on *why* it is correct or appropriate to do so, on what *makes it correct*. It seems to me that, at least in many cases, what makes it correct to value something for its own sake is that the thing is independently valuable—that is, its value comes from its nature.

Of course, from a general perspective, the hierarchical structure of practical reason implies the explanatory priority of our more general ends. If, to use Garthoff's (2010) example, Scrabble expertise constitutes having a hobby, and having a hobby constitutes using and improving my cognitive capacities, then it seems obvious that the value of using and improving my cognitive capacities will import its value down, so to speak, into the specific aim of Scrabble expertise. Since I clearly have a reason to use and improve my cognitive capacities, I therefore have a reason to adopt Scrabble expertise as an end. But there are many other ways to have a hobby and to use and improve my cognitive capacities (some of them presumably better than Scrabble, at least in some ways and for some agents). Thus, the line of explanatory dependence must run from using and improving cognitive capacities (the more general end) to Scrabble expertise (the more particular end that potentially constitutes the general one). But I think that reflection on the qualitative experience of being drawn to an activity by its particular

idiosyncrasies suggests that a complete explanation will actually run in both directions. After all, the universal must be instantiated in the particular.

Here it will be helpful to return to Brewer's accounts of dialectical activity and desire, and to Wolf's account of meaningfulness in life.

Some of Brewer's central examples of dialectical activities are activities of friendship and love. These kinds of activity are valuable because other persons have independent worth. Thus, the value of these activities has a source that is independent of the role those activities play (or might potentially play) in constituting anyone's good life, well-being, or flourishing. For these kinds of activities, it is easy to see how our understanding of a completely independent value (the worth of another person) drives our understanding of the value of the activity. Activities of friendship and love are partly constituted by an earnest appreciation of the independent worth of the other person. Activities of friendship and love get their value from the nature of their constitutive objects.

If Brewer's account of dialectical activity and evaluative outlooks is correct, then paradigmatically an agent is drawn into a friendship through an apprehension of the other person's worth. This is the value toward which one's desires direct one's attention when one desires the friendship of a particular person. Of course, friendship in general is desirable for its own sake, and also as a constitutive element of a meaningful life. So one might be tempted to say that the general value of friendship, of having *a* friend, explains the value of this particular friendship, at least for the individual pursuing the friendship. Moreover, the value of having a meaningful life explains the value of having a friend. But I think reflection on the worth of individual persons warrants the conclusion that the value of a particular friendship is in some ways more fundamental than the value of friendship in general. Reasons of my well-being and

flourishing are not the only reasons to which I ought to be responsive; there are also other sentient and rational beings as well, and they have worth that merits various forms of appreciation, regardless of whether that appreciation contributes to my own good (Garthoff 2010: 709, fn. 35). I can have reasons to pursue a friendship that have nothing to do with how that friendship will contribute to the goodness of my life, and I think that these reasons are the most fundamental reasons for understanding what makes friendship valuable. When I am drawn to a person in this way, considerations of my own good or even of my friend's good are hardly operative at all, at least in the (admittedly somewhat idealized) paradigmatic case. It is not that these considerations have no bearing; rather, they are overshadowed and unnecessary to account for what makes a friendship choiceworthy. Interestingly, from this point of view it seems that the value of having *a* friend depends upon the way that the general activity of friendship draws our attention to and makes practically significant various idiosyncrasies of individual persons, and thus brings us into proper contact with individual persons who have independent worth. This worth does not depend at all on how relations with individual persons figure in anyone's well-being, or the goodness of anyone's life. Rather, the reverse is true: a person's good matters because of the worth that the individual person has.

I think full appreciation of the worth of an individual person requires a grasp of the peculiar qualities of that individual. There are, of course, impersonal ways to appreciate and respect the worth of persons, even persons we know nothing about. And these impersonal responses to individual worth are important, since we cannot (and should not) have close personal relationships with everyone. But the impersonal appreciation of the worth of persons is, I think, less fundamental than the kind of appreciation that is essential to more intimate

friendships. Thus, I think reflection on how we grasp the worth of persons in the case of close relationships is essential to understanding the nature of that worth.¹⁵⁰

Turning our attention toward meaningfulness can help bring this point into view. Meaningfulness in one's life is not the same as well-being. Any living creature has a well-being, but meaningfulness seems to require rational nature, and perhaps *critically* rational nature. Yet meaningfulness is not just the flourishing or well-being of critically rational individuals. The claim "It is good for you, as a critically rational being, to have a meaningful life" is not a tautology. We cannot develop an understanding of the notion of a meaningful life without a prior understanding of particular meaningful activities and projects. And I do not think this limitation is just a quirk of our psychology. The reason that we can reasonably disagree about whether having a meaningful life is necessarily good for a person is that the notion of meaningfulness encompasses a kind of value that is not essentially teleological or relative. That is, the species of value that meaningful activities and projects bring into view is not just goodness relative to function, flourishing, or well-being. The most obvious example of this kind of value is the worth that individual persons have.

As Wolf emphasizes, the reasons one has to pursue meaningful activities are not essentially reasons of self-interest (2012: 1–7). She calls the reasons we have to pursue meaning in our lives "reasons of love" (4).¹⁵¹ This phrasing suggests an important qualification of the view. When agents act out of reasons of love, they are responding to value in a way that, in the

¹⁵⁰ Brewer (2021) argues that it is love, not respect, that most fully articulates the "irreplaceable value" of persons. His point is that respect, although an important and often appropriate response to the worth of persons, cannot be the ground of morality's authority.

¹⁵¹ Wolf self-consciously borrows the phrase "reasons of love" from Harry Frankfurt (2004), though I think she develops the idea in a more promising direction than Frankfurt does. I find the phrase useful, but I think that there may be reasons that play the same kind of role in meaning-making that are not best captured by the term "love," such as aesthetic reasons.

right conditions, brings meaning to their lives. But they are not pursuing meaningfulness directly. They are not, that is, acting under the guise of the meaningful, and the reasons they do act on—reasons of love—are entirely self-sufficient and would continue to be motivating even if an agent somehow concluded (erroneously) that these activities do not contribute to her life’s meaningfulness (117–118).¹⁵²

I think that these considerations show the fundamentality of the particular in explanations of meaningfulness. In the context of a discussion of human well-being, it makes sense to explain the value of an activity or project in terms of how that activity or project contributes to the good of the individual pursuing it. Since individuals can and often do act out of direct concern for their own good, attending to goods like well-being and flourishing brings the hierarchical structure of practical reason into view. With this perspective, one values an end like Scrabble expertise insofar as it constitutes a more general end, like having a hobby, all the way down to the most general end, that of living a good life. But this kind of explanation can make it seem quite puzzling that in order for a particular end like Scrabble expertise to help constitute more general ends like living a good life, one must value the particular end for its own sake. Why does valuing Scrabble expertise *merely* for the sake of having a hobby or the sake of living a good life undermine the power of that end to constitute those more general ends? If we widen our view beyond the scope of well-being and flourishing, I think we can begin to see the shape of the solution.

¹⁵² Wolf claims that moral reasons are not like this, since an agent moved by moral reasons might cease to be so moved if she were to conclude that morality did not actually demand she act on those reasons (117). There is something to this argument, but I think she overstates it. I think that moral reasons often have the same structure as the reasons of love that Wolf focuses on. That is, particular moral reasons—like reasons of kindness—need not present themselves under the guise of the moral in order to be motivating. Particular moral reasons may be more sensitive to the general guise of the moral, however, since the agent *should* cease to be moved by them if she sincerely concludes that morality demands that she *not* act on them. How much reasons of morality and reasons of meaning really differ in this respect deserves further consideration.

The value of meaningfulness does not fully explain the value of particular meaningful activities. Rather, the value of the general end, meaningfulness, is itself partly explained by the value of the particular ends that constitute it. In other words, a meaningful life is choiceworthy *because of* the value of the meaningful activities, projects, and relationships that constitute it. And while a full explanation of the value of those particular activities, projects, and relationships may require some reference to how they contribute to meaningfulness (as well as other goods like well-being), the general goods are not the most fundamental values overall. These considerations lead me to conclude that meaningfulness *just is* the value that an individual's life has when it is lived in close and proper relation to those things that have independent worth—to what matters. This explains why in order to achieve meaningfulness, one cannot aim at meaningfulness directly. Rather, in order to live a meaningful life, one must pursue the independently choiceworthy objects and activities that constitute meaningfulness because that is what constitutes the proper relation between the individual and what matters. And insofar as meaningfulness is the excellence toward which practical thought should strive, we can understand practical thought as constitutively functioning to bring the individual into proper relation with what matters.

And finally, now we can see why meaningfulness occupies a special place in the hierarchy of practical reason, even though it is not the most comprehensive good. The good life is a more general good than the meaningful life, since meaningfulness partly constitutes the good life. But I think meaningfulness is nonetheless the most important general good around which practical reason ought to be organized. This is because meaningfulness is the value that an individual's life has when the individual lives in proper relation to the things that have independent worth—most paradigmatically, other persons and sentient beings.

IV. Conclusion

I conclude with some brief reflections on the choiceworthiness of human life in general. I began this chapter with the observation that the question of what makes human life itself worthy of choice is one that ethical theory must wrestle with in one way or another. Christine Korsgaard (1986) argues that this question demands an answer because of the structure of justification:

Justification, like explanation, seems to give rise to an indefinite regress: for any reason offered, we can always ask why. If complete justification of an end is to be possible, something must bring this regress to a stop; there must be something about which it is impossible to ask why. This will be something unconditionally good. Since what is unconditionally good will serve as the condition of the value of other things, it will be the source of value. (488)

Korsgaard's point, which she makes in the context of comparing the (meta)ethics of Aristotle and Kant, is that there are many things that are choiceworthy *given that one is a human being*. But these ends are only choiceworthy under the contingent conditions of human life. She finds in both Aristotle and Kant the idea that for our pursuit of these conditional goods to be fully justified, there must be at least one *unconditional* good, and that this unconditional good will be what makes human life itself choiceworthy (502–505). It will be the source of the value of other goods.

If my arguments here are on the right track, then there is something amiss about this formulation of the problem. If I am right that the value of particular activities, projects, and relationships helps explain the meaningfulness of a human life, then it is these same particular ends—worthwhile activities, projects, and relationships—that make human life choiceworthy. These particular ends are “conditionally good” in the sense that their value may not be intelligible outside of a human (or at least rational) point of view. But at least in some cases, like

friendships, the source of the value of the end (the friendship) is the unconditional (that is, completely independent) worth of the end's object (the individual friend).¹⁵³ It seems clear that the worth of human persons is unconditional, but it is equally clear that human persons are not the source of their own worth—at least not in Korsgaard's constructivist sense, where the ultimate explanation for why human persons have worth is that we must suppose that they do in order for our rational choices to be intelligible. Even if it is true that the intelligibility of our rational choices depends upon our treating human persons as unconditionally valuable, that at best offers only an indirect reason to think that human persons in fact have a special kind of value (which I call independent worth). It does not explain the source or ground of that value at all. Moreover, human persons are probably not the only things with worth. Non-human conscious animals, for example, seem to me also to have independent worth.

These points are warranted by reflection on the nature of evaluative representation. The nature of representational kinds depends on the possibility of (at least approximate) veridicality. Any philosophical paradigm that supposes a category of representations that are veridical merely by virtue of being formed in the right way, with no reference to the way the world independently is, creates for itself a substantial burden of argument. Nor is it any better to deny, as non-cognitivists do, that evaluations are representational. If we keep a clear view of the phenomenological, cognitive, and perceptual resources necessary for the complex network of evaluative representations that compose practical thought, then I think the evaluative outlook view surfaces as the most plausible framework for thinking about practical reasoning and the way that we engage with values in our lives.

¹⁵³ More investigation is needed into the difference between those activities, like friendship, that center on objects of obvious worth, such as other persons, and those, like hobbies, that have a less direct connection to independently worthwhile things.

References

- Begby, E. (2011) “*Origins of Objectivity*” (book review), *Notre Dame Philosophical Reviews*.
- Birch, J. (2022) “The Search for Invertebrate Consciousness.” *Noûs*, Vol. 56, No. 1, pp. 133–153.
- Block, N. (2002) “Concepts of Consciousness.” In David J. Chalmers (ed.), *Philosophy of Mind: Contemporary Readings*. Oxford University Press, Oxford, pp. 206–218.
- Brady, M. (2013) *Emotional Insight: The epistemic role of emotional experience*. Oxford University Press, Oxford.
- Brewer, T. (2009) *The Retrieval of Ethics*. Oxford University Press, Oxford.
- . (2021) “Acknowledging Others.” *The Journal of Ethical Reflections*. Vol. 1, No. 4, pp. 91–115.
- Burge, T. (1995) “Content Preservation,” *Philosophical Issues*, Vol. 6, pp. 271–300.
- . (2003) “Perceptual Entitlement,” *Philosophy and Phenomenological Research*, Vol. 67, No 3, pp. 503–548.
- . (2007a) “Two Kinds of Consciousness.” in Tyler Burge, *Foundations of Mind*, Oxford.
- . (2007b) “Reflections on Two Kinds of Consciousness.” in Tyler Burge, *Foundations of Mind*, Oxford.
- . (2009) “Five Theses on *De Re* States and Attitudes.” In J. Almog and P. Leonardi (Eds.), *The Philosophy of David Kaplan*, Oxford University Press, Oxford.
- . (2010a) *Origins of Objectivity*. Oxford University Press, Oxford.
- . (2010b) “Origins of Perception.” *Disputatio*, Vol. IV, No. 29, pp. 1–38.
- . (2010c) “Steps Towards Origins of Propositional Thought.” *Disputatio*, Vol. IV, No. 29, pp. 39–67.

- . (2014) “Reply to Block: Adaptation and the upper border of perception,” *Philosophy and Phenomenological Research*.
- . (2022) *Perception: First form of mind*. Oxford University Press, Oxford.
- Campbell, J. (2011) “Tyler Burge: Origins of Objectivity.” *Journal of Philosophy* 108, pp. 269–285.
- Carey, S. (2009) *The Origins of Concepts*. Oxford University Press, Oxford.
- Cordner, C. (2014) “Dialectical Activity, Ritual, and Value: A critique of Talbot Brewer.” *Philosophical Investigations* Vol. 39, No. 2, pp. 178-191.
- Crane, T. and C. French. (2021) "The Problem of Perception." *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/fall2021/entries/perception-problem/>](https://plato.stanford.edu/archives/fall2021/entries/perception-problem/).
- Davidson, D. (1980) *Essays on Actions and Events*. Oxford University Press, Oxford.
- Deonna, J. (2006) “Emotion, Perception and Perspective.” *Dialectica* Vol. 60, No. 1: pp. 29–46.
- Donnellan, K. (1966) “Reference and Definite Descriptions.” *The Philosophical Review*, Vol. 75, No. 3, pp. 281–304.
- Döring, S. (2007) “Seeing What to Do: Affective perception and rational motivation.” *Dialectica*, Vol. 61, No. 3, pp. 363–394.
- Elgin, C. (2007) “Understanding and the Facts,” *Philosophical Studies*, Vol. 132, No. 1, pp. 33–42.
- . (2008) “Emotion and Understanding.” In G. Brun, U. Dogouglu, and D. Kuenzle (eds.), *Epistemology and Emotions*, Ashgate Publishing, Farnham.
- Frankfurt, H. (2004) *Reasons of Love*. Princeton University Press, Princeton, NJ.

- Freas C. A., Spetch M. L. (2023) "Varieties of visual navigation in insects." *Animal Cognition* 26 (1): 319-342. doi: 10.1007/s10071-022-01720-7. Epub 2022 Nov 28. Erratum in: *Animal Cognition* 2022 Dec. 13: PMID: 36441435; PMCID: PMC9877076.
- Frege, G. (1892) "Über Sinn und Bedeutung."
- Garthoff, J. (2010) "Structuring Ends." *Philosophia* 38, pp. 691–713.
- Grimm, S. (2006) "Is Understanding a Species of Knowledge?" *British Journal of the Sciences* Vol. 57, No. 3: pp. 515–535.
- . (2016) "Understanding and Transparency." In *Explaining Understanding: New perspectives from epistemology and philosophy of science*, Routledge.
- Higgenbotham, J. (1998) "Conceptual Competence." *Philosophical Issues* 9: pp. 149–162.
- Hills, A. (2009) "Moral Testimony and Moral Epistemology." *Ethics* Vol. 120, No. 1: 94–127.
- Kant, I. (2002) *Groundwork for the Metaphysics of Morals*. T. Hill (ed.) and A. Zweig (trans. & ed.), Oxford University Press, Oxford.
- Korsgaard, C. (1983) "Two Distinctions in Goodness." *The Philosophical Review*. Vol. 92, No. 2, pp. 169–195.
- . (1986) "Aristotle and Kant on the Source of Value." *Ethics*, Vol. 96, No. 3, pp. 486–505.
- . (1996) *The Sources of Normativity*. Cambridge University Press, Cambridge, MA.
- . (1997) "The Normativity of Instrumental Reason." In G. Cullity & B. N. Gaut (eds.), *Ethics and Practical Reason*. Oxford University Press, New York.
- . (2003) "Realism and Constructivism in Twentieth-Century Moral Philosophy." *Journal of Philosophical Research* 28, pp. 99–122.
- . (2008) "From Duty and for the Sake of the Noble." in *The Constitution of Agency*, Oxford University Press, Oxford, pp. 174–206.

- . (2018) *Fellow Creatures: Our obligations to other animals*. Oxford University Press, Oxford.
- Kim, J. (1994) “Explanatory Knowledge and Metaphysical Dependence”, *Philosophical Issues* 5: pp. 51–69. Reprinted in Kim 2010: 167–186. doi:10.2307/1522873.
- . (2005) *Physicalism, Or Something Near Enough*. Princeton University Press, Princeton, NJ.
- Marcus, E. (2012) *Rational Causation*, Harvard University Press, Cambridge, MA.
- Moore, G. E. (1922) “The Conception of Intrinsic Value.” *Philosophical Studies*, Rutledge and Kegan Paul, Ltd.
- Nagel, T. (1974) “What It Is Like To Be a Bat.” *Philosophical Review* 83, pp. 435–450.
- . (1986) *The View from Nowhere*. Oxford University Press, Oxford.
- Noë, A. (2004) *Action in Perception*. MIT Press, Cambridge, MA.
- Prinz, J. (2004) *Gut Reactions: A perceptual theory of emotions*. Oxford University Press, Oxford.
- Richardson, H. (1997) *Practical Reasoning About Final Ends*. Cambridge University Press, Cambridge.
- . (2004) “Thinking About Conflicts of Desire.” In P. Baumann and M. Betzler (eds.), *Practical Conflicts: New philosophical essays*. Cambridge University Press, Cambridge, MA.
- Scanlon, T. (1998) *What We Owe to Each Other*. Harvard University Press, Cambridge, MA.
- Shapiro, T. (2014) “What Are Theories of Desire Theories Of?” *Analytic Philosophy*, Vol. 55, No. 2, pp. 131-150.

- Siegel, S. (2021) "The Contents of Perception." *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/fall2021/entries/perception-contents/>](https://plato.stanford.edu/archives/fall2021/entries/perception-contents/).
- Stampe, D. (1987) "The Authority of Desire." *Philosophical Review*, Vol. 96, No. 2, pp. 335–382.
- Tappolet, C. (2017) *Emotions, Value, and Agency*. Oxford University Press, Oxford.
- Tenenbaum, S. (2007) *Appearances of the Good: An essay on the nature of practical reason*. Cambridge University Press, New York.
- Tumulty, M. (2020) *Alien Experience*. Oxford University Press, Oxford.
- Wolf, S. (2012) *Meaning in Life and Why It Matters*. Princeton University Press, Princeton, NJ.
- Wiggins, D. (1987) "A Sensible Subjectivism." In Wiggins, D. *Needs, Values, Truth: Essays in the philosophy of value*. Oxford University Press, Oxford.
- Williams, B. (1979) "Internal and External Reasons." In Ross Harrison (ed.), *Rational action: studies in philosophy and social science*. New York: Cambridge University Press. pp. 101-113.
- Wringe, B. (2015) "The Contents of Perception and the Contents of Emotion." *Noûs* Vol. 49, No. 2: pp. 275–297.
- Zagzebski, L. (2003) "Emotion and Moral Judgment." *Philosophy and Phenomenological Research* 66: pp. 104–124.

Vita

Michael Ebling grew up in Gainesville, Florida. He moved to Tennessee to pursue his undergraduate degree at Sewanee: The University of the South. After receiving his BA in philosophy from Sewanee, he completed an MA in teaching secondary English from Vanderbilt University, and he spent several years teaching 7th and 8th grade English in public schools. He then decided to pursue a doctorate in Philosophy, which led him to study at the University of Tennessee in Knoxville. Besides philosophy, Michael is especially interested in rock climbing, game design, and music.