ICONS OF VALUE: A STUDY OF THE EPISTEMIC SIGNIFICANCE OF EMOTIONAL
EXPERIENCE

by

Brian Scott Ballard

BA, University of California Santa Cruz, 2009

MA, New York University, 2011

Submitted to the Graduate Faculty of

The Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of Doctor of Philosophy

University of Pittsburgh

2017

UNIVERSITY OF PITTSBURGH

THE DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Brian Scott Ballard

It was defended on

May 4, 2017

and approved by

Karl Schafer, Professor of Philosophy, University of California Irvine

Crispin Wright, Professor of Philosophy, New York University

James Woodward, Distinguished Professor of History and Philosophy of Science,
University of Pittsburgh

John McDowell, University Professor of Philosophy, University of Pittsburgh

Dissertation Director: Edouard Machery, Distinguished Professor and Director for the Center for
Philosophy of Science, University of Pittsburgh

ICONS OF VALUE: A STUDY OF THE EPISTEMIC SIGNIFICANCE OF
EMOTIONAL EXPERIENCE

Brian Ballard, PhD

University of Pittsburgh, 2017

The relation between emotion and value is one of the great longstanding issues in value

theory. Anti-realists often have ready accounts of the nature of this relation. For instance,

some anti-realists hold that emotions somehow construct evaluative facts. However, it is

less obvious how a realist should understand the emotion-value relation. Some realists

leave emotions out of the picture altogether, while others regard them as pernicious

influences. The present work seeks to address this. My goal is to characterize the

emotion-value relation in a way that is amenable to a robust form of value realism, and

which gives emotions an indispensible epistemic role. The key theses I argue for are (1)

that emotions provide epistemic benefits with respect to value but that presently no

account succeeds in saying how they do this and what those benefits are; (2) that

emotions represent evaluative properties, and that each emotion represents a unique

evaluative property; and (3) that emotions represent evaluative properties in a distinctive

way that can be modeled after pictorial styles of representation.

TABLE OF CONTENTS

LIST OF FIGURES

PREFACE

I first started to think about emotions as an undergraduate at UC Santa Cruz in 2008. Jesse Prinz visited campus for a talk, and I was fortunate enough to have been able to arrange a coffee, where we discussed emotions and metaethics. Back then I had a clear idea that I wanted a shamelessly robust realism that took seriously the importance of emotional experience. This I struggled to convey to Jesse, who, with his equally shameless anti-realism, described it as a "neat disgusting project." I was very much drawn to two ideas then, which I carry with me still. The first is that we know about values through a kind of direct experience. The second is that this evaluative experience, whatever it is, is not metaphysically neutral about the values it conveys. It somehow conveys them as being mind-independent. I thought then, and think now, that this is a feature of sense perception as well. Samuel Johnson once kicked a rock and claimed to have refuted Berkeley. But perhaps his refutation isn't as feeble as it seems. Perhaps the point is that, when you kick the rock, this experience makes it feel as though you were encountering an object that exists independently of your mind. I wanted to be able to say something similar about the experience of value. In encountering the beauties of nature or the horrors of a war, it feels as though one were kicking the proverbial rock.

I thought it would be powerful to combine these ideas with a third one. G. E. Moore and Thomas Reid both took an approach to skepticism that I found captivating.

They thought it so obvious that we have perceptual knowledge, that any philosophical premises implying otherwise must be treated in the form of a modus tollens. I liked the idea that our everyday experiences enjoy a powerful default warrant that philosophical considerations do not have the strength to overturn. (As an undergraduate, I had the good fortune of discussing these Moorean-Reidian ideas with John McDowell and Crispin Wright, who, I note with pleasure, sit on the committee of this dissertation presently.) Now, the way that this third idea combines with the first two is obvious. Our everyday experiences acquaint us with values which we encounter as mind-independent, and so philosophical doubts about the mind-independent reality of value must perish along with skeptical doubts about perceptual knowledge.

These were the ideas I felt in my guts, that dark place where philosophy begins. But it's quite another thing to draw them out into the light in a way that will convince. This dissertation attempts a very modest beginning.

My work on these ideas has been developed throughout many conversations and correspondences, and I wish to thank those who have been so generous with their time and wisdom, including Mikio Akagi, Bob Audi, Tim Bayne, Salim Berker, Bill Brewer, Dan Brudney, Alex Byrne, Agnes Callard, Elizabeth Camp, Dave Chalmers, Robert Cowan, Jennifer Corns, Tim Crane, John Deigh, Kit Fine, Matt Frise, Ronald de Sousa, Alan Gibbard, Camil Golub, Anil Gupta, Joshua Greene, Stephen Grimm, Chris Hill, Paul Horwich, Mark Johnston, Thomas Kelly, Chris Korsgaard, Joseph LeDoux, Barry Maguire, Joe McCaffrey, Michael Milona, Jerry Neu, Richard Nisbett, Graham Oddie, Derek Parfit, Adam Pautz, Jesse Prinz, Duncan Pritchard, Jim Pryor, Hilary Putnam, Robert Roberts, Tim Scanlon, Kieran Setiya, James Shaw, Daphne Simeon, Jay Van

1.0 INTRODUCTION

1.1 THE IMPORTANCE OF EMOTIONAL EXPERIENCE

We can understand the importance of emotions by considering what life is like without them. In his work on autism, Oliver Sacks writes about his trip through the Rockies with Temple Grandin, the autistic scientist. The following passages are worth quoting at length:

> We pulled off the road, and gazed toward the Rockies—snowcapped, outlined against the horizon, luminously clear even though they were nearly a hundred miles away. I asked Temple if she did not feel a sense of their sublimity. "They're pretty, yes. 'Sublime,' I don't know." When I pressed her, she said that she was puzzled by such words, and had spent much time with a dictionary, trying to understand them. She had looked up "sublime," "mysterious," "awe," and "numinous," but they all seemed to be defined in terms of one another.

> "The mountains are pretty," she repeated, "but they don't give me a special feeling, the feeling you seem to enjoy." After living for three and a half years in Fort Collins, she said, this was only the second time she had been to them.

Later, as they watched the sunset, Grandin remarked:

> "You get such joy out of the sunset," she said. "I wish I did, too. I know it's beautiful, but I don't 'get' it." Her father, she added, often expressed similar sentiments.

And again, while walking under the night sky:

>"When I look up at the stars at night, I know I should get a 'numinous' feeling, but I don't. I would like to get it. I can understand it intellectually. I think about the Big Bang, and the origin of the universe, and why we are here: Is it finite, or does it go on forever?"
>
>"But do you get a *feeling* of its grandeur?" I asked.
>
>"I *intellectually understand* its grandeur," she replied, and continued, "Who are we? Is death the end? There must be reordering forces in the universe. Is it just a black hole?" (Sacks 1995, pp. 293-294)

Clearly, Grandin is missing something with respect to the aesthetic quality of the mountains, the sunset, and the night sky. Aesthetics is not entirely alien to her, as she can recognize that these things are 'pretty' and 'beautiful'. Still, she seems to be, in a sense, blind to the more dramatic aesthetic qualities of grandeur and sublimity, and her 'blindness' here seems to be at least partly explained by her lack of emotional experiences such as awe and wonder. The Grandin example makes plausible, then, at least the very generic claim that emotional experience is *in some way* crucial to judgments of aesthetics.

In fact, this goes beyond aesthetics. Imagine the moral analogue of Grandin:[1]

>Sacks: Can you believe those hoodlums were about to poor gasoline on the poor cat? Isn't that just horrible?

---

[1] It is sometimes argued that psychopaths are real-life examples of this (Prinz 2008, pp. 42-47), although how to understand the nature and significance of emotion deficits in psychopathy is a complicated matter, as is no doubt the case with autism as well.

Moral Grandin: 'Horrible', I don't know. It was certainly wrong and they deserve to be punished.

Sacks: But don't you get a *feeling* of the wrongness?

Moral Grandin: I *intellectually understand* the wrongness. We ought to treat animals with respect, and it is bad to enjoy the pain of another.

It is natural to say something similar about this case. If one is not distressed by the hoodlums' act, then one does not fully grasp its badness. It is when you feel appalled that you really 'get it'. Plausibly, then, we may expand our generic claim: emotional experience is in some way crucial to judgments of *value*.

In what way exactly? Just what is Grandin missing on account of her emotional deficits? This is the guiding question of my dissertation. By way of preview, my central claims will include the following: (1) that emotions are representations of value; (2) that emotions are not beliefs or judgment but are instead experiences that are largely encapsulated from beliefs, much like perception is; (3) that as evaluative experiences, emotions furnish epistemic benefits that purely intellectual states cannot; (4) that current work on emotion is unable to explain why emotional experience is so epistemically valuable; (5) that the thing that sets emotion apart from intellection is that emotions represent iconically, by replicating aspects of the values they represent; and (6) that emotions are iconic in virtue of their distinctive phenomenal character.

In the rest of this introduction, I will describe the issues that form the landscape of these ideas, and end by offering a preview of each of the three chapters of the dissertation, which can be read independently of each other.

1.2 EMOTION AND VALUE THEORY

Many metaethical theories emphasize that there is a deep connection between emotion and value, and offer ready explanations of what Grandin is missing. For *non-cognitivists*, what Grandin is missing is the emotional life that forms the real substance of evaluative judgment, and so she lacks genuine evaluative judgments with respect to the values in these specific cases (Blackburn 1993, 1998; Gibbard 1990, 2003). Of course, on this view, there are no evaluative properties, and our evaluative judgments are not trying to describe them, not in anything more than a deflationary or minimalist sense; and so it is more apt to say that her evaluative *judgment* is deficient, rather than saying her grasp of an evaluative property is. (Of course, quasi-realist views make things complicated here, since such views are willing to talk of evaluative properties and even say that they are mind-independent.)[2]

In contrast, *sentimentalist theories* accept that there are evaluative properties, but hold that they are in some way dependent on our emotional responses. This is true of *response-dependence theorists*, who think that values can be fully explained in terms of the emotional responses they tend to cause (Hume 1738-40/1975, 1742/2006; Hutcheson 1724/2004; Johnston 1989; Lewis 1989; Prinz 2008). It is also true of *fitting attitude*

---

[2] On the increasing unrecognizability of non-cognitivism as an alternative to realism, see Dreier (2004).

theorists, who think that values can be fully explained in terms of the emotional responses they make fitting (D'Arms 2013; D'Arms and Jacobson 2000, 2006; McDowell 1985, 1998a, 1998b; Wiggins 1976, 1987). Note, however, that the way D'Arms and Jacobson develop this position should be distinguished from the way McDowell and Wiggins develop it. We can see the difference by considering how they answer the Euthyphro question, which Crispin Wright has shown to be a useful way of contrasting metaphysical positions in the vicinity here.[3] We can all agree that an act is pious *iff* loved by the gods. The question is whether the gods' love is *detecting* piety or instead *constructing* it. In the present case, D'Arms and Jacobson hold that emotions are metaphysically prior to values, and so emotions are in some way constructing values. In contrast, McDowell and Wiggins hold that emotions and values are metaphysically or explanatorily symmetrical.[4] The beautiful is that which merits admiration, and admiration is that which the beautiful merits, a circularity which proponents of this view embrace. To keep this distinction in mind, I will call McDowell and Wiggins' approach *the sensibility theory*, and reserve the label 'fitting attitude theory' for D'Arms and Jacobson's approach. The sensibility theory is more realist in spirit, especially in McDowell's presentation, although it is a subtle issue just how to classify sentimentalist theories in relation to realism (Dancy 1986; Pettit 1991; Tollefsen 2000).

Sentimentalist theories appeal to an analogy with color, taking it for granted that color must be understood as a secondary quality. Indeed, for some reason, just about all of metaethics seems to regard this as the only possible view about color. That is odd,

---

[3] For details, see Wright (1987), pp. 42-46, (1988), pp. 14-25, and especially (1993), pp. 108-139.
[4] In response to Blackburn's distinction between realist views on which values are "the parents of our sentiments", and anti-realist views on which they are "their children", McDowell remarks that values and sentiments are *siblings* (McDowell 1998b, p. 159). See Blackburn (1981), pp. 164-165.

since there are as many views about color as there are about value.[5] Consider, for example, Thomas Reid:

> All people who have not been tutored by modern philosophy understand by color, not a sensation of the mind, which can have no existence when it is not perceived, but a quality or modification of bodies, which continues to be the same whether it is seen or not (Reid 1872, p. 137).[6]

Here, the idea is that the reality of color is metaphysically and explanatorily prior to our color experiences. This is accordingly a deeper form of realism than the secondary quality view of color, since it makes color even more independent of our minds. And like sentimentalist theories of value, I, too, wish to appeal to an analogy with color, only I understand color as Reid describes it, and not on the secondary quality model. Accordingly, one of the aims of my dissertation is to show how a robust form of realism about value can explain what Grandin is missing: *Values are metaphysically prior to emotion, but emotion is a way of detecting values, and it is a way of detecting values that is epistemically better than non-emotional ways of doing so*.

I think of values as being like color in yet a further respect: their manifold diversity and determinacy. Just as visual experience makes subtle discriminations among a vast range of highly determinate colors, so does emotional experience make subtle discriminations among a vast range of highly determinate values (Johnston 2001). In

---

[5] For instance, Parfit writes, "Though some such *response-dependence* theory must be correct when applied to colours, … there are strong objections to response-dependent accounts of morality" (2011, p. 379). However, a huge volume of contemporary work on color rejects response-dependence accounts. There we find a diverse range of realist views, including reductive views (Byrne and Hilbert 2003; Hilbert 1987; Jackson 1998, 2007; Matthen 1988; Tye 2000) as well as non-reductive primitivist views (Allen 2010, 2017; Campbell 1994, 2005; Gert 2008; Hacker 1987; McGinn 1996; Watkins 2002).

[6] In fact, although tempting, the question of what is there when no one is looking is the wrong question for contrasting primary and secondary quality views of color, since an object can have the *disposition* to cause certain color experiences, even when hidden away at the bottom of the sea.

emotional experience, we encounter the austere majesty of the Alps, the terrible beauty of the Pacific, the sacred beauty of wild horses, the cuteness of puppies, the crisp goodness of early autumn in New York, the creepiness of abandoned houses, the bitter injustice of police brutality in the US, the other-worldly evil of the things serial killers do to their victims. This list is long but not nearly long enough. Every value admits of a vast range of determinates. The canyons of Utah and the sea cliffs of California are both sublime, but each in its own way, and someone who can't see this would make a lousy novelist. Moreover, these determinates are reflected in emotional phenomenology. The fear one feels walking past a graveyard is not the same as the fear one feels having stumbled upon a rattlesnake. This is partly a phenomenal difference. We emotionally sense the distinct ways we are vulnerable in each scenario.

To be a realist about so many values is to offend against the preference for parsimony that prevails in much of philosophy today. This picture of a world haunted with values is certainly not Quine's taste for desert landscapes. No doubt, parsimony considerations have been brought against value realism even of a less extravagant kind (e.g., Harman 1977, Ch. 1). It is, however, difficult to say just what parsimony is, let alone what its value is (Sober 2001, 2015). Indeed, much of history has had the "taste" opposite to Quine's, expecting every nook and cranny of the universe to be filled with something (Lewis 1964/2012, p. 56; Lovejoy 1936). A related 'principle of plenitude' has sometimes been appealed to in physics (Dirac 1930, p. 71, n5), and continues to be discussed in metaphysics (Hawthorne 2006; Inman 2014; Leslie 2011). So, we should not be naively optimistic about the damage parsimony can do to the lavish realism I endorse.

I submit, moreover, that my picture at least takes emotional experience at face value, as many have observed. For instance, McDowell remarks, "Aesthetic experience typically presents itself, at least in part, as a confrontation with value: an awareness of value as something residing in an object and available to be encountered" (1998a, p. 112).[7] What to do with this observation is partly a matter of one's approach to philosophy. While some are happy to cast aside such appearances as illusory (Mackie 1977), I'll note that many regard them as furnishing a powerful default justification (Kelly 2005; Lemos 2004; Lycan 2001, 2007; Moore 2013, essays VII, IX, and X; Reid 1872, Part A, Ch. I).

At any rate, in my dissertation, although I will not attempt an all-out defense of realism, my project is a kind of defense, albeit indirectly. This is so in two respects. First, one of the great advantages of sentimentalist theories of value is their ready account of the deep connection between emotion and value.[8] Realists have a harder time achieving such an account, and many realists do not seem concerned to try (e.g., Huemer 2005; Nagel 1986; Parfit 2011). Yet that is just what my dissertation seeks to provide. Thus, if I am successful, it will dull the anti-realist's edge, or at least, this particular edge. I wish to show, in other words, how even a lavish form of realism can give emotional experience a central place with respect to value and judgments of value.

Second, I will argue that emotions are representations of value, and while some non-cognitivists may be able to *say* this, the way that I get to this claim is not compatible with a minimalist or quasi-realist treatment of representation. What this claim ends up being, then, is in a sense the deepest way of rejecting non-cognitivism. Not only do we

---

[7] McDowell, of course, does not think this objective phenomenology requires us to be realists in any sense more robust than the sensibility theory he adopts.

[8] This advantage is emphasized, for instance, in Blackburn (1998), Ch. 4.

have evaluative beliefs, but even our emotions are evaluative representations. Our evaluative psychology is representational all the way down. Just imagine Ayer's chagrin if evaluative utterances expressed feelings, yet feelings themselves turned out to have evaluative content.

1.3 EMOTIONS AS SOURCES OF EVALUATIVE JUDGMENTS

If emotional experience is an epistemically important way of detecting values, then it must be that emotional experience is *upstream* of our evaluative judgments. One line of emotion research, however, regards emotional experience as taking place downstream. You decide what is beautiful, and your longings merely reflect this prior judgment. You decide what is unjust, and your indignations merely reflect this prior judgment. If this is right, then emotions have a less interesting epistemic role to play (Bagnoli 2011; Sinnott-Armstrong 1991). For instance, plausibly they could not serve to justify the evaluative beliefs they reflect, since that would seem to license a kind of objectionable bootstrapping—the beliefs cause the emotions, which justify the beliefs.[9]

In the case of aesthetics, however, this picture is manifestly implausible. A child does not need to have a prior conception of aesthetic value in order to be moved by a sunset or a cello sonata, anymore than one needs a prior conception of culinary value in order to enjoy a glass of water. Of course, one needs the conceptual sophistication to understand certain works of art before one can be moved by them. But such concepts are

---

[9] An analogue of this issue is found in older debates about the "theory ladenness of observation". If observation (perception) is as theory laden as some have argued (Kuhn 1964, p. 113-114), then this would limit its role as a source of information about the world, as critics have noted (Fodor 1983, 1984). A similar issue arises in more recent debates about the epistemic impact of the cognitive penetrability of perception (Siegel 2017, Ch. 3).

not *evaluative*. The concepts needed to understand Shakespeare are not BEAUTY and SUBMLIMITY, but IAMBIC PENTAMETER and TRAGIC THEATRE. And that is unsurprising. In order for emotion to evaluate something, we must have *some* sense of what it is (although apparently that "sense" can be as minimal as subcortical processing early in perceptual systems (Zajonc 1980; LeDoux 1996, Ch. 6)).

The relevant question, then, is the extent to which emotions depend on prior *evaluative* cognitions, since it is evaluative cognition for which I am claiming emotional experience to be epistemically important. To be clear, one of the background assumptions of the present work is that emotions play a substantial role in influencing evaluative judgment, that often we judge something to be valuable or disvaluable partly as a result of emotional experience. This is plausible in aesthetics, and there is significant evidence that it is plausible in ethics as well (Haidt 2001, 2012).[10] For instance, in one of Jonathan Haidt's famous studies, participants who sat at disgusting desks were likely to pass more severe moral judgments (Schnall et al. 2008). Importantly, this effect was modulated by the extent to which subjects were sensitive to their bodily feelings, which suggests that conscious affect is playing a role. Beyond ethics, furthermore, emotions seem to influence cognitions of other values as well, such as judgments of overall wellbeing (Clore et al. 2001; Schwartz and Clore 1983).

The idea that emotions are sources of evaluative judgments, rather than mere reflections of them, also explains two further facts, namely, that (a) emotional experiences may lead us to *revise* our evaluative judgments, and (b) emotional experiences may *conflict* with our evaluative judgments. Accordingly, given that

---

[10] But see Huebner et al. (2009) for a contrary take on this evidence.

emotions influence evaluative judgments as antecedents, plausibly they have a more robust epistemic role to play. At the very least, it raises the question of whether the influence of emotion is something we should welcome.

1.4 SKEPTICISM ABOUT EMOTION

Certainly, there are those who think we should not welcome it. They regard emotion as a distorting influence likely to lead us away from evaluative truth. The many disparate traditions of such *emotion skepticism*, as we may call it, form an impressive roster.

Most (in)famously, there are the stoics. As Cicero remarks, "all emotion springs from the roots of error: they should not be pruned and clipped here and there, but yanked out" (2002, p. 60). Similarly, Kant often discusses emotions in a way that views them as interruptions in the moral life, rather than central features of it:

> The true strength of virtue is a *tranquil mind* with a considered and firm resolution to put the law of virtue into practice. That is the state of health in the moral life, whereas an affect, even one aroused by the thought of what is good, is a momentary sparkling phenomenon that leaves one exhausted (1996, p. 516).[11]

More recently, there are three contemporary brands of emotion skepticism worth mentioning. The first are those ethical intuitionists who regard cognitive intuitions as the primary source of ethical knowledge, and emotion as a distorting influence (Huemer 2008). The second are those who regard emotion's allegedly more prominent role in deontological judgments as a debunking argument against deontological approaches to

---

[11] Interpreting Kant's actual view of emotion is doubtless a more complicated affair, one that requires discussion of his aesthetics as well.

ethics (Greene 2003, 2008; Singer 2005, p. 347-348).[12] The third are researchers in the heuristics and biases tradition who have, it seems, discovered patterned ways the emotions can mislead us (Kahneman 2011; Nisbett and Ross 1980; Slovic 2007).

How can we overcome such a venerable community of emotion skeptics? Obviously, I cannot systematically address all of this work, but there are general considerations to bring against emotion skepticism. For starters, the Grandin case from earlier is relevant. Without being awed by the night sky, Grandin seems unable to fully register its aesthetic value, even though she has an intellectual understanding of its grandeur. We should regard emotions as epistemically beneficial in at least this respect.

Moreover, it is plausible that we have some kind of default warrant to trust in our natural ways of forming beliefs. This approach is advocated by proponents of bounded rationality (Gigerenzer and Selter 2001; Simon 1957), as well as many epistemologists (Alston 1989, 1993; Conee and Feldman 2004, Ch. 1; Plantinga 1967; Wright 1991, 2004). If that is correct, then we do not need to *earn* a warrant for believing emotions are reliable. Rather, provided that emotions are one of our basic sources of evaluative judgments, we start out getting to trust in them.

Of course, such a default warrant would be defeasible, and some emotion skeptics are confident there are defeaters. For instance, Slovic (2007) exhibits evidence that emotions are most sensitive to information that is concrete and vivid, and so pallid statistics about, say, mass genocides tend not to be registered in the emotions. However, as an argument for emotion skepticism, this makes the mistake of confusing sensitivity with reliability. Here, a belief source is insensitive *iff* it fails to represent the information

---

[12] See Berker (2009) for sustained criticism of the relevance of Greene's findings for any normative conclusions.

that is there; unreliable *iff* it misrepresents the information that is there.[13] The blind person's vision is insensitive, the hallucinating person's unreliable. With this distinction in view, it is clear that the mere fact that emotions only produce verdicts in response to certain kinds of data doesn't show that the verdicts they *do* produce are unreliable.

Indeed, even setting that point aside, the claim that emotions are unreliable is intrinsically problematic. It assumes there is such a thing as the reliability of emotion, as though all people had the same emotions. We can speak of the reliability of perception in a global way since most humans tend to have relevantly similar perceptual faculties. But the same cannot be said for emotions. Emotions depend heavily on training and refinement to a far greater degree than perception. There is no such thing as the global reliability of emotion, any more than there is such a thing as the global morality of the will. The real question is: When emotions have been adequately trained and refined, what, then, is their epistemic impact? And as the Grandin case suggests, when emotions get things right, their epistemic impact is indispensible.

1.5 THEORY OF EMOTION

If emotions are to play an important role in value theory, we must address, What *are* emotions? Here is an oversimplified story of emotion theory in the last hundred years. First came the feeling theory. On this view, emotions are just certain kinds of bodily feelings (Damasio 1994; James 1884; Lange 1885). This view respected the salient role of feeling in emotion, but, it was thought, had difficulty accommodating other things

---

[13] For a more precise characterization of the sensitivity-reliability distinction, see Sober (1994), pp. 51-52.

about emotions, such as their intentionality, and the way they can be assessed as fitting. Then there came the cognitive theory. This view hoped to fix these problems by saying that emotions are evaluative beliefs or judgments (Foot 1978a, 1978b; Kenny 1963; Nussbaum 2001; Solomon 1976). After all, judgments are intentional. And judgments can be true or false, so perhaps the fittingness of emotion is just a kind of truth or falsehood. This theory, however, was vulnerable to obvious counterexamples from so-called recalcitrant emotions. You can fear the insect even when you judge it to be harmless. Thus, there came the perceptual theory, which holds that emotions are perceptions of value (Döring 2003; Prinz 2004; Roberts 1988, 2003; Tappolet 2012). This theory seemed to have it all. Perceptual states are intentional. Perceptual states can be assessed for accuracy, and this can explain the fittingness of emotion. Moreover, perceptual states can come apart from judgments. Thus, perhaps recalcitrant emotions are like visual illusions, persisting even as they conflict with our better judgments.[14]

There are of course many other theories of emotion, but this story is a useful oversimplification. And the perceptual theory is still the view to beat, although at present it faces a challenge that it has not yet entirely come to terms with. Emotions admit of normative assessments that do not comfortably apply to paradigmatic perceptual states (Brady 2007, 2009, 2011, 2014, Ch. 3; Helm 2001, 2015). For instance, emotions can be irrational, but it is awkward to say that a visual state, even a hallucination, is irrational (but for a contrary view, see Siegel 2017). This is currently the main problem to

---

[14] What about moods that apparently lack intentional objects? I think we should compare these to ganzfeld experiences in which the visual field is pervaded with a single uniform color. On one view, the ganzfeld experience represents a *plane* as colored. On another view it represents an *unattributed* color property. Both views are possible with respect to moods. Perhaps depression represents *things* as bad, or perhaps it represents unattributed badness. See Mendelovici (2014) for a more detailed view.

overcome for the perceptual theory, and naturally, several solutions have been put forth (Döring 2015; Milona 2016; Tappolet 2012; Vance forthcoming).

In the present work, while I proceed in a way that is cozy with the perceptual theory, nothing I say depends on the claim that emotion is literally a form of perception. Instead, the claims I rely on are (a) that emotional experiences are largely encapsulated from evaluative beliefs, (b) that emotional experiences are often causally and epistemically prior to evaluative beliefs, (c) that one of the functions of emotion is to register values, (d) that emotional experience *presents* its evaluative content as being correct, and (e) that emotional experience is a source of epistemic goods with respect to values. Although these claims highlight perceptual features of emotion, they do not strictly entail that emotions are perceptual states.

1.6 PREVIEW OF CHAPTERS

Here I offer a brief summary of the three chapters.

1.6.1 CHAPTER ONE: EMOTIONS AS SOURCES OF EPISTEMIC GOODS

I argue that emotional experience furnishes epistemic benefits with respect to value, but that current work on emotion is unable to fully explain this. I argue this by offering a series of evaluative Mary cases, in which someone shares our evaluative judgments but forms them without emotion. Like Grandin, such a person seems to be epistemically worse off with respect to value. But how exactly? I consider whether this can be

explained in terms of attention, perception, or understanding, and find that none of these notions are entirely adequate.

1.6.2 CHAPTER TWO: EMOTIONS AS REPRESENTATIONS OF VALUE

Many authors have claimed that the fittingness of emotion directly supports the view that emotions have evaluative content. However, this inference is mistaken, since *actions* can be fitting, but they do not have evaluative content. In this chapter I fix the inference by showing how the fittingness of emotions can figure in a more complex argument for the evaluative content of emotion. I argue that each emotion has a unique set of fittingness conditions *and* a unique phenomenal character. All things being equal, it would be preferable to have an explanation that *unifies* these two features, rather than having two unrelated explanations of each. But the best unifying explanations, I argue, commit us to the claim that emotional experience represents evaluative properties.

1.6.3 CHAPTER THREE: EMOTIONS AS ICONS OF VALUE

I develop a novel semantics for the evaluative content of emotion, and I do so in three steps. First, I introduce the notion of *arbitrariness*. Here, a representation is arbitrary to the extent that we can change its format without changing its meaning. English words are highly arbitrary in this sense, but realistic pictures are not. Second, I argue that emotional experience is non-arbitrary. If we alter certain aspects of emotional phenomenology, such as its pleasantness or painfulness, we alter its evaluative content. Third, I introduce the notion of *iconicity*, and argue that this notion can help us explain why emotional experience is non-arbitrary with respect to value. Here, a representation is iconic to the

extent that it represents in virtue of replicating certain properties of its object. Plausibly,

iconic representations are highly non-arbitrary, and so if emotions are icons of value, this

would explain their non-arbitrariness.

2.0 EMOTIONS AS SOURCES OF EPISTEMIC GOODS

There are two kinds of value realists—those who think emotions give us a kind of epistemic access to the evaluative facts, and those who do not. The former view we may call **sentimentalist realism**. This view has gained a serious following in recent years (Audi 2013; Brady 2014; Cowan 2015; Cuneo 2006; Dancy 2014; Deonna and Teroni 2012; Döring 2003; Johnston 2001; Milona 2016; Oddie 2005; Pelser 2014; Roberts 2013; Tolhurst 1991; Wedgwood 2007). For instance, Mark Johnston regards emotional experience as a way of sensing "patterns of appeal and repulsiveness there in the environment itself" (2001, p. 186). Similarly, Robert Roberts remarks that, when an injustice is observed,

> one who feels the injustice for himself, by way of his indignation, has an epistemically *higher-quality* judgment than the emotionless person. The [emotional] experience of the injustice gives him deeper understanding and more intimate cognitive contact with this moral reality (2013, p. 52).

Notice that as Roberts puts things, emotional experience is both necessary and sufficient for the relevant epistemic good. The idea is that there is a *distinctive* epistemic benefit arising from emotional experience, not attainable otherwise. There is, of course, a weaker version of sentimentalist realism, which holds that while emotional experience provides an epistemic good, it can be attained in some other way as well, perhaps through cognitive intuition (Audi 2013; Cowan 2015; Dancy 2014; Wedgwood 2007).

Sentimentalist realism contrasts with sentimentalist non-realism on which emotions construct values, rather than disclosing them (D'Arms and Jacobson 2000, 2006; Lewis 1989; Mulligan 1998; Prinz 2008). Both views agree that emotions and values have an important explanatory connection, but they disagree about whether it is emotion or value that has the explanatory priority. In this chapter, I will set non-realist views aside, and focus entirely on realist views. My aim is to address how realists should conceive of the relation between value and emotion, given that realists hold values to be in no way metaphysically dependent on emotion (or any other mental state). This is relevant to the dialectic between realists and non-realists, since non-realists often claim to have a better account of the emotion-value relation (e.g., Blackburn 1998).

The kind of realism I am addressing here regards value as metaphysically prior to emotion and other human responses. This excludes some views that count as realist on other taxonomies, such as the sensibility theories of McDowell (1985) and Wiggins (1976). On these views, values are understood in terms of the emotions they merit, and emotions are understood in terms of the values that merit them, a circularity which defenders of the view regard as benign. These views can be construed as realist in a certain sense: Sunsets can be objectively beautiful, because it can be an objective fact that they merit admiration. But these views can be construed as non-realist with respect to how they answer the Euthephro question: neither values nor emotions have metaphysical priority, but both are explained in terms of the other. Accordingly, in this chapter, I will regard such views as non-realist, since I am interested in stronger forms of realism on which values have metaphysical priority over emotions. It is with respect to these realist views that the problem arises of explaining the relation between value and emotion.

Not all realists think that emotions give us a kind of epistemic access to value. This is especially so with respect to *moral* realism. Indeed, across the full spectrum of moral realist views, emotions are often left out entirely (Boyd 1988; Brink 1989; Enoch 2011; Jackson 1998; Nagel 1986; Parfit 2011; Scanlon 2014; Shafer-Landau 2003). Other moral realists argue positively that emotions cannot play certain significant epistemic roles, for instance, as sources of foundational justification for moral beliefs (Brogaard and Chudnoff 2016). Still others regard emotions as positively *hindering* epistemic access to ethical truths (Huemer 2008).[15] Let **non-sentimentalist realism** include both the view that emotions hinder epistemic access to value, as well as the view that they simply play no significant epistemic role with respect to value.

Sentimentalist realists can admit that emotions sometimes hinder our epistemic access to value, and non-sentimentalists can admit that emotions sometimes help it. The question is, rather, the extent to which emotions have a general, favorable role in the epistemology of value. The non-sentimentalist will say this takes place at most in isolated cases.

In this chapter, I want to address both kinds of realism. To the non-sentimentalist realists, I will argue that they are mistaken: emotional experience provides *some* manner of epistemic benefit with respect to value, and does so with regularity. To the sentimentalist realists, I will argue that extant accounts fail to capture just what the epistemic benefit of emotional experience is.

---

[15] Parfit has also communicated this view to me in conversation, remarking that—and I paraphrase here—"if emotion influences moral belief to the degree the moral psychologists say, then I would regard that as a debunking explanation of moral belief."

In section 2.1, I argue that emotional experience provides *some* manner of epistemic benefit. What exactly is this epistemic benefit, and why does emotional experience provide it? The rest of this chapter addresses potential answers. In section 2.2, I consider the view that emotions have important attentional effects. In section 2.3, I consider the view that emotions are perceptual experiences of value that provide immediate justification for evaluative beliefs. In section 2.4, I consider the view that emotions provide evaluative understanding, where understanding is distinct from justification and knowledge.

While this chapter is largely negative, emphasizing the ways that extant theories fall short, this negative work is useful at present, given the proliferation of views concerning the epistemic significance of emotion. Note, however, that this chapter is not entirely negative. I will argue for the positive thesis that emotional experience provides some kind of epistemic good. While this thesis is often either taken for granted (Brady 2014, Ch. 1), or dismissed out of hand (Huemer 2008), it is important to see that it can be argued for. I turn to that task now.

2.1 THE BENEFITS OF EMOTIONAL EXPERIENCE

In this section, I will present a series of what I call evaluative Mary cases. In these cases, it is intuitive that there is *some* benefit being provided by emotional experience. I will argue that the benefit cannot be fully explained as a *moral* benefit, and so we should think it is *epistemic*.

2.1.1 EVALUATIVE MARY CASES

Imagine that you are hiking in the mountains with your friend, Dud. You can see a meadow through the trees ahead. As you approach, you realize that you are not alone: In the meadow is grazing a herd of wild horses. You watch from the tree line unnoticed, their manes rustling in the wind, the strength of their haunches defined in the midday light. What do you feel when you see this?

I would feel something like reverent awe, the sort of longing that is also a kind of pain. One might also feel simple joy, or amazement, or gratitude if one believes in God. I assume you, too, would feel something in this vicinity.

Dud, however, seems to feel nothing at all. He fidgets with his compass. "Dud," you say, "are you seeing this?".

"Yeah."

"Isn't it amazing?"

"Yeah."

But his response, cold and aloof, leaves you vaguely unsatisfied. You want him to feel the thrill that you feel, to encounter these wild horses as you encounter them.

If you were to describe your dissatisfaction, to say just what is wrong with Dud's response, you might mention a number of things. The point is first to notice that there is *something* Dud is missing, *some* sense in which you are better off than Dud, although I say nothing yet about what that is.

This is an evaluative Mary case because it depicts the difference between experientially responding to a certain property, and merely knowing about that property; then the case invites us to explain just what the difference is.[16]

It is worth seeing that such cases can be constructed for a range of values. Here are three further cases.

Baby example: You and Dud finally see your mutual friend's newborn baby. As you watch the baby, she gives you and Dud a big smile. You feel that familiar cooing feeling. Dud, however, feels nothing, but he agrees the baby is cute.

Thin ice example: You and Dud are snow shoeing together, and you both slip and fall onto a frozen lake. Naturally, you feel terrified. Dud, however, feels nothing, but he agrees the situation is dangerous.[17]

---

[16] There have been two other explicit attempts to construct an evaluative analogue of the Mary case from Jackson (1982), each being restricted to moral value. First, in Prinz (2008), a moral Mary case is used to argue that certain emotional experiences (disgust, shame, etc.) are the sources of our ordinary moral concepts. I've set up my evaluative Mary case to neutralize this verdict, since Dud is not generally devoid of emotional experiences, and so he possesses whatever concepts they might provide. Second, in Yetter-Chappell and Yetter-Chappell (2013), a moral Mary case is used to argue for an unbridgeable gap between moral and natural concepts. Their case does not explore the role of emotion. In contrast, my evaluative Mary case is set up deliberately to assess the role of emotional experience in value epistemology, just as one might use the original Mary case to assess the role of visual experience in color epistemology, an obvious application of Jackson's thought experiment that has not yet been fully considered. It is nearly always assumed that Mary gains something regarding the look of red—but why not the red itself? If you're a direct realist about perception, that just may be the thing to say.

[17] This is a variant of a case from Stocker (1983).

Genie example: You're reading a book about the Genie Clark case, the worst recorded case of child abuse in US history. Horrified, you read passages out loud to Dud about how she was kept in a dark room and bound to her potty chair. You feel sickened and saddened. Dud, however, feels nothing, but he agrees it's terrible what was done to Genie.

I take it to be an obvious datum in each of these cases that there is *some* respect in which Dud is worse off than you are. I say nothing yet about how exactly he is worse off, only that he is. This is the intuitive starting point of my argument. If this much does not seem correct to you, if it seems instead that there is no way in which you are better off than Dud, then I have no argument to convince you otherwise.

I assume, moreover, that in these cases Dud's agreement expresses an evaluative belief. He really does think the wild horses are beautiful, etc. Some anti-realists will balk at this, rejecting the possibility of emotionless evaluative belief. I have said nothing to refute such anti-realist positions, as they are not my targets here. Instead, this chapter addresses how realists should assess the epistemic profile of emotional experience, and realists generally allow for the possibility of emotionless evaluative belief. Of course, a sentimentalist realist could in principle hold that occurrent emotional experience is required for evaluative belief. But I'm aware of no realists who actually hold this strong view.

Barring anti-realist views, we seem to be able to make evaluative judgments without consciously feeling emotion. For instance, is it worse to kill 100,000 people, or 1,000,000? Of course the latter. But when we consider crimes of that magnitude in terms

of pallid data, we often feel unmoved (Slovic 2007). So, I take it the Dud cases are familiar enough. After all, it isn't that Dud is *globally* without emotional experience. He's just gone cold in these specific situations.

Plausibly, the way you're better off than Dud arises from emotional experience, since that is the main difference between the cases. The question is, What exactly is the benefit you gain? *Sentimentalist* realists, of course, will say the benefit is in some way epistemic. But non-sentimentalists will have to explain it in some other way. After all, if in these four cases emotional experience provides an epistemic benefit, then plausibly we should think this about many other cases as well.

Since these are evaluative Mary cases, one might think the most popular view about the original Mary case from Jackson (1982) will be applicable here, namely, that you and Dud know the same fact in different ways (Churchland 1985; Horgan 1984; Levine 2007; Loar 1990; Tye 1986). One might even think that's *all* there is to it: neither of you know the fact better, only differently. Now, that may be, but in that case, this view wouldn't explain the sense in which you are better off than Dud for having been emotionally moved. And that is just what we are trying to explain at present. If this view is to explain the sense in which you are better off than Dud, then it must be the case that the way you know the fact is somehow better than the way Dud knows it. This, however, would be to understand the benefit of emotional experience as epistemic. And that is just what I am trying to convince you of at present.

Let us begin by addressing non-sentimentalist realists who want to characterize the benefit as non-epistemic. If the benefit is not epistemic, what is it exactly? The most natural answer is that it is some kind of moral benefit. I'll discuss the most natural versions of this view, and argue that none of them work.

2.1.2 MOTIVATION

Perhaps you are better off than Dud because you are more motivated to behave in the appropriate ways. This fits especially well with the baby and the thin ice examples. In these cases, there are behavioral responses that are important for you to be ready to carry out, say, caring for the baby, or escaping from the ice.

But this fits awkwardly with the wild horses example. Sure, some motivation is provided by your emotional experience, but to do *what*? Perhaps to linger, or take a picture, or write a poem; but unlike the baby and thin ice examples, there is no behavior we think you really must do, such that the emotional experience gets its importance as something instrumental in getting you to do it. Granted, there are actions which your feelings towards the wild horses motivate you *not* to do, such as harming them, or even just scaring them off for fun. But we needn't suppose you were in danger of doing these things, were it not for your feelings of longing.

The same is true of the Genie case. There is nothing you need to do, since the thing has already happened. Sure, perhaps the motivated behavior is a future one, say, to give your own children a good life, or to care for the vulnerable and oppressed. Good things to do, no doubt, but even if your being moved by the Genie case *didn't* lead to any change in your life projects, it still seems better to have been moved by the story.

2.1.3 FITTINGNESS

Perhaps the difference is simply that you have a fitting emotion and Dud lacks one. However, on most accounts, the fittingness of emotion arises from the *evaluative content* of emotion (Döring 2003; Greenspan 1988; Helm 2001; Kenny 1963; Neu 1977; Prinz 2004; Roberts 2003; Solomon 1976; Tappolet 2012). It is fitting to be angry at injustice because anger *represents* injustice. Of course, some sentimentalist *non*-realists might resist this (D'Arms and Jacobson 2000; Gibbard 1990), since this makes values prior to emotions. But these are not the targeted views at present.

Given that fittingness comes from evaluative representation, the fittingness of your emotion can't be the whole story. After all, Dud has a correct evaluative representation—his evaluative belief. So, it must be that there is something special about the *way* emotional experience represents values. This may indeed be a good way to go, but notice that now it isn't mere fittingness that's doing the work.

Notice as well that, if we go this route, then plausibly we will end up describing an *epistemic* good, since we will be describing a good that arises in virtue of a distinctive kind of representation. It will be an epistemic good at least in the sense that accurate representation is an epistemic good.

2.1.4 VIRTUE

Perhaps the difference is that Dud lacks certain virtues. For instance, maybe your being saddened by the Genie case reflects your compassion. But this can't be the right story. Dud might have the relevant virtue as well, even if it doesn't manifest on this occasion

for some accidental reason. More plausibly, the difference would be that Dud doesn't feel what the virtuous person would feel *qua* virtuous person.

It is certainly true he doesn't. But what the virtuous person would feel is *the fitting emotion*. And it is plausible that the virtuous person feels it *because* it is the fitting emotion. This is especially plausible, given that fittingness arises from evaluative representation. Consider an analogy with virtue epistemology. Perhaps the virtuous cognizer would tend to hold only true beliefs. But that is not what *makes* the beliefs true.[18] Now watch the explanation play out: Why would the virtuous person be saddened by the Genie case? Because sadness is fitting. Why is sadness fitting? Because sadness accurately represents it as bad. But if sadness represents it as bad, then plausibly what we have here an epistemic good, that is, the good of accurate representation. We end up with the same dialectic we had with the appeal to fittingness.

2.1.5 CONCERN

Perhaps the problem is with Dud's *concerns*. In the wild horses example, perhaps he is not adequately concerned with natural beauty, even though he is able to recognize it. However, as with virtue, Dud might possess the relevant concern, if it simply fails to manifest on this occasion for some accidental reason. Similarly, *you* might be unconcerned with natural beauty. Perhaps being moved by the wild horses is a fleeting isolated incident for you. Even so described, the intuitive betterness of your having the emotional experience doesn't go away. We don't think the emotional experience is better simply in so far as it reflects your prior concerns.

---

[18] But for a contrary view, see Blackburn (2001).

2.1.6 MORALITY

More generally, it is implausible that Dud is only morally worse off than you are, given the range of cases. It isn't natural to think of his lack of emotion in the wild horses example, for instance, as subject to moral assessments, such as blame and punishment, and we don't think your being moved by the wild horses is morally praiseworthy. It isn't a matter of *duty*, for instance. Same with the thin ice example. Assuming you and Dud both try to escape, and thus fulfill any duty to preserve your lives, it doesn't seem that you are morally *admirable* for having been afraid, or that Dud is morally *reprehensible* for having not been. This sort of language doesn't fit.

2.1.7 FINISHING THE ARGUMENT

To recap, the non-sentimentalist realist rejects the claim that emotional experience provides any special epistemic access to values. But in the evaluative Mary cases, it is intuitive that your emotional experience provides *some* benefit. So, the non-sentimentalist will have to explain it some other way. The most natural answer is that it is a moral benefit, but we have seen this cannot entirely explain the difference between you and Dud.

If we were to sum up the argument, it would look like this:

1. The benefit from your emotional experience is either moral or epistemic.

2. It is not moral.

3. So, it is epistemic.

I've argued for 2, but what about 1? Admittedly there are other kinds of benefits besides moral and epistemic ones. However, we can quickly see they fail to provide the explanation we need. For instance, perhaps the benefit is *pragmatic*. In the ice example, perhaps you are more likely to escape given your fear. Sure, but this is does not generalize, since there doesn't seem to be any pragmatic benefit in being moved by the wild horses. Indeed, if I offered a million dollars to the person who is unmoved by the wild horses, Dud would be pragmatically better off, but this does not erase the other sense in which he is worse off than you.

Alternatively, you might appeal to norms of biological function. Perhaps Dud's emotions are not functioning right, and that's what's worse about him. However, it seems too strong to say that Dud is unmoved because his emotions are *malfunctioning*. That could be why, but needn't be.

In that case, perhaps emotional experience provides a *sui generis* benefit. But we should regard this as a last resort. If the benefit can be understood as epistemic, then we should resist postulating a *sui generis* normative status.

Are there problems with the idea that the benefit of emotional experience is epistemic? One might object here that emotions are *downstream* of the epistemically important stuff. When you experience the right emotions in the evaluative Mary cases, this merely reflects your *prior* evaluative understanding. So, you do have better epistemic access to values than Dud has, but your emotions merely evidence this rather than generating it.

However, it isn't correct that emotional experience only ever reflects one's prior evaluative understanding. Consider *evaluative conversions*. Imagine someone who is racist towards group X. But one day, when he sees a member of X in abject misery, he feels moved with pity, and this feeling brings home to him the humanity of X. If this leads him to abandon his racist attitudes, then his emotional experience would have altered his evaluative conception, generating new evaluative understanding.

Of course, there may be other objections worth considering, and we can't address them all here. Still, I take it a case has been made for the verdict that what you gain over Dud is epistemic. This verdict also resonates with the fact that we can describe evaluative Mary cases in terms that sound overtly epistemic. We can say that you *better grasp* or *more fully register* the beauty of the wild horses, for instance. This sort of language seems to fit.

Assuming this is right, what exactly is the epistemic benefit furnished by emotional experience? Realists can appeal to a variety of answers here. I'll consider three current accounts, and argue they do not fully explain the epistemic benefits of emotional experience.

2.2 ATTENTIONAL EFFECTS

Some theories of emotion emphasize its *attentional impact*. For instance, Ronald de Sousa argues that emotions are "determinate patterns of salience among objects of attention, lines of inquiry, and inferential strategies" (1987, p. 196), and are thus capable of "guiding the processes of reasoning—or distorting them…" (p. 197). It is in this sense that "emotions set the agenda for beliefs and desires" (p. 196).

More recently, Deonna and Teroni have emphasized this feature of emotion (see also Brady 2014, Ch. 5). They claim that "[g]iven the complexity of the environment through which we navigate, the prospects for detecting properties that could justify evaluative judgments without the aid of [emotional] sensitivity are not promising" (2012, p. 121). The idea here seems to be that not only do emotions direct attention, but that lacking emotions would seriously damage our ability to attend to the features relevant to making good evaluative judgments.

Assuming this is correct, perhaps this explains what is wrong with Dud. Perhaps the epistemic benefit of emotional experience is that it makes us better able to attend to the evaluatively relevant features of the scenario.

However, we should reject this explanation. It may be true that having the right emotions improves our ability to attend to the relevant features. But it is clearly false that we *must* have the right emotions in order to so attend. In the evaluative Mary cases, then, I stipulate that even though Dud is unemotional, he nevertheless does succeed in attending to the relevant features of the scenario—the apparent strength of the wild horses, their breath misting in the cool afternoon. He attends to these things, albeit coldly,

and it is on the basis of so attending that he forms the correct evaluative beliefs. Even so, this does not seem to erase the sense in which you are better off for feeling moved.

Alternatively, perhaps it is a question not of success but ability. Perhaps it is that your *ability* to attend to the relevant features is better. However, we needn't imagine the evaluative Mary cases in this way. Just suppose that you suffer from attention deficit, while Dud does not. The mere reduction in your attentional abilities does not seem to lessen the sense in which you are better off for having been moved by the wild horses.

I conclude, then, that while emotions clearly have attentional effects, and while this is clearly one important aspect of their epistemic significance, it cannot fully explain the epistemic good that emotional experience provides.

2.3 PERCEPTUAL JUSTIFICATION

Recent philosophical work on emotion has been dominated by the so-called perceptual theory of emotion, the view that emotion is a form of perception. This suggests that the perceptual features of emotion can explain the epistemic difference between you and Dud.

The version of the perceptual theory we will consider holds that emotions are perceptual states with *evaluative* content (Döring 2003; Prinz 2004; Roberts 1988, 2003; Tappolet 2012). This view is preceded by the cognitive theory of emotion, which holds that emotions are evaluative *judgments* (Foot 1978; Nussbaum 2001; Solomon 1976). When you are afraid of something, you *judge* it to be threatening. However, many reject the cognitive theory in light of the possibility of recalcitrant emotions (see Roberts 2003,

pp. 83-105).[19] You might fear the garden snake even though you know it is harmless. In such cases, far from being an evaluative judgment, your fear seems to *conflict* with your evaluative judgment. Instead, it is attractive to think of such cases as akin to known perceptual illusions. The mirage still looks like water, even when you judge it to be a mirage.

If emotions are evaluative perceptions, this is obviously of interest to the epistemology of value, and that has not gone unnoticed. As Sabine Döring writes:

> The fact that emotions have representational content opens up the possibility that the occurrence of an emotion can, in suitable circumstances, entitle a thinker to judge, and possibly to know, its content simply by taking its representational content at face value (2003, p. 229).

As evaluative perceptions, it is natural to think emotions can justify evaluative beliefs. This epistemological upshot of the perceptual theory was noted early on as an attraction of the view. Roberts (1988) was one of the first versions of perceptual theory, and the epistemological significance of his view was noticed soon after by Tollhurst (1991). Since then, many have defended the idea that emotions justify evaluative beliefs on the model of perception (Cowan 2015; Cuneo 2006; Döring 2003; Johnston 2001; Milona 2016; Pelser 2014; Roberts 2013). Of course, independently of the perceptual theory of emotion, virtue ethicists promoted the idea of moral perception (e.g., McDowell 1979), and accordingly, virtue ethics and the perceptual theory of emotion have a natural kinship (e.g., Roberts 2013). Indeed, Aristotle himself may have held a kind of perceptual theory of emotion (Dow 2011), and the same can be argued for Aquinas (King 2002).

---

[19] There are other reasons to reject the cognitive theory as well; see Deigh (1994) and Griffiths (1997).

The idea, then, is that emotions are evaluative perceptions, and that *this* can explain the epistemic difference between you and Dud. However, I want to show that this idea is problematic. Indeed, I issue here a general explanatory challenge for the perceptual theory: the theory in its extant versions is unable to fully explain the epistemic significance of emotional experience.

The problem is this. Emotion may be a form of evaluative perception, but there may be other forms of evaluative perception as well, or perhaps ways of forming evaluative beliefs that share all the epistemically relevant features of perception. Some of these might be entirely non-emotional. So, if the epistemic benefits of emotional experience derive merely from its perceptual features, then non-emotional forms of evaluative perception will in principle provide the same exact epistemic benefits. But this is a problem for two reasons. First, many proponents of the perceptual theory hold that emotions provide epistemic benefits that are not otherwise attainable. Recall the Roberts passage from the start of this chapter, in which he claims the evaluative judgment of "the emotionless person" is epistemically worse. Second, as I will argue in a moment, merely intellectual states do not provide the epistemic benefit that emotions do, even when the intellectual states share the perceptual features of emotion.

Sense perception justifies in virtue of certain of its features rather than others. For instance, my visual experience does not justify in virtue of the fact that my eyes are above my nose. What, then, are the justification-making features of perception, the features in virtue of which sense perception provides the justification it does? Many perceptual theorists of emotion are neutral on that point, and assume that whatever the story is in the case of sense perception, it will transfer to emotion, provided that emotions

are perceptions. As Döring puts it, how it is emotions justify will be explained "by analogy to the entitlement in sense perception" (2003, p. 229). Other perceptual theorists, however, are more explicit about what they take to be the justification-making features of perception. Using these more explicit views, I'll develop my challenge to the perceptual theory.

2.3.1 EMOTIONS AS RELIABLE INDICATORS OF VALUE

Cuneo (2006) adopts a kind of indicator reliabilism about the justification we get from emotions. Cuneo uses the word "feelings" instead of "emotions", but he has in mind "feelings of aversion, disgust, repulsion, uneasiness, distress, guilt, attraction, pleasure, relief, delight and the like" (2006, p. 70). So, he seems to have in mind emotions as well as states with similar affective phenomenology. Cuneo holds that in many cases, such feelings

> play a special sort of role in generating … value judgments. … If we assume that in cases such as these, feelings reliably indicate value, then they also function as *signs* or indicators of evaluative features of the world (2006, p. 69-70).

For Cuneo, then, feelings justify because they are reliable signs of value. (Of course, sometimes they are unreliable.)

Notice, however, that just about anything can in principle be a sign of value. According to Livy, a chicken refusing to eat is a sign of ill things to come (1912, X 40). We can imagine a world where such things are reliable omens. To use a more grounded example, many philosophers hold the view that cognitive intuitions, seemings, or appearances justify a range of moral and other beliefs (Bealer 1998; Huemer 2005;

Chudnoff 2013).[20] These intuitions are usually taken to be states that can be experienced without any accompanying affect. And many proponents of the perceptual theory of emotion are happy to grant this. They are willing to say that both intuition and emotion are sources of evaluative justification.

But in that case, let us imagine that Dud forms his evaluative beliefs on the basis of cold intuition. Let us stipulate, moreover, that Dud's cold intuitions are as reliable as your emotions. Both you and Dud are using reliable signs of value. However, this does not erase the sense in which Dud is missing something. Just imagine that as you watch the horses, you and Dud continue your dialogue:

"Dud, do you find these horses amazing or what?"

"Yeah."

"Well you seem unimpressed."

"Look, I'm a bit tired. But I can tell they are beautiful."

"How can you tell?"

"Well, they just seem beautiful."

Even when their beauty is intellectually apparent to Dud, you somehow more deeply appreciate their beauty when you are moved by it. This means that, even if Cuneo is right, and emotions justify as signs of value, this doesn't fully explain the epistemic benefits of emotional experience, since intellectual intuitions can also be reliable signs of value without conferring the same benefits. Now, if there were something special about

---

[20] A different notion of intuitions regards them as beliefs that are somehow self-justifying (Audi 2005).

the *kind* of signs that emotions are, things could be different. But that of course requires further theorizing.

At this point, some perceptual theorists of emotion may want to deny the existence of cold intuition, and go in for the strong view that emotion is the *only* form of evaluative perception. But this view faces serious challenges from cases in which we seem to make unemotional evaluative judgments in an immediate way. For one, consider judgments about permissibility: when you judge an act to be permissible, this is often enough an immediate moral judgment formed without emotion (Dancy 2014, p. 796). Moreover, consider judgments about distant situations. Imagine that you could stop a careening trolley in a distant country by pressing a button on your computer that would drop a human being on the tracks. Should you? It seems to me No, but I don't feel especially worked up about it. Concreteness, physicality, proximity, particularity—these factors hold great sway over our emotions, but our immediate moral judgments can often break free from them.

Further, even if you deny that our *actual* psychology is equipped with cold intuitions, it's enough if such intuitions are *possible*. After all, I just need Dud to have intellectual seemings in my imaginary examples. So, if the intuition skeptic wants to resist my argument, she has to go in for the stronger view that intuitions are not even possible.

2.3.2 EMOTIONS AS APPEARANCES OF VALUE

So much for reliable indication. Many perceptual theorists of emotion, however, think that emotions provide justification in virtue of being appearances or seemings, regardless of their reliability. Many think this about emotions or affective experiences, even if they do not identify as a perceptual theorist of emotion. For instance, Jonathan Dancy defends the claim that emotions are "practical seemings" or "presentations of reasons" (2014, p. 796). Here, the key features of seemings are the following:

Belief independence: It can seem to you that p, even if you believe not-p.

Representation: Seemings have representational content, and can therefore be assessed in truth terms (e.g., as true or false, accurate or inaccurate).

Presentation: Seemings *present* their representational content as how things really are. This is a phenomenal quality. When you shut your eyes and imagine a pair of hands, it doesn't feel as though the hands are really there. But when you open your eyes, and see your hands, it feels as though the hands are real. Your visual experience *presents* hands.

Some authors describe additional features of seemings, but these are the key features for our purposes. The notion of presentation is especially important. For those who wish to eschew externalist conditions on justification, presentation is often cited as the justification-making feature of perception. Huemer calls this feature "assertiveness"

(2001, p. 53), and Pryor "phenomenal force" (2000, n37), but they are referring to the same aspect of perceptual phenomenology, and they both consider it to be the justification-making feature of perception. Plausibly, this phenomenal feature is also present in other states as well, such as cognitive intuitions and some memories.

As noted, many perceptual theorists hold that, since emotions are evaluative perceptions, the justification-making feature of perception, whatever it is, will be had by emotions as well. Thus, if Huemer and Pryor are right about the justification-making feature of perception, then it wouldn't be enough to say, as Döring does, that emotions represent evaluative properties; emotions would also need to *present* their evaluative content.

As with Cuneo's view, the problem here is that it is plausible we can also have purely cognitive intuitions. I've motivated this view briefly in 3.1, and anyways, many will find it unattractive to resist the view, even those who defend the epistemic importance of emotion. Now, intuitions—on the view relevant for our purposes—are understood to be appearances or seemings that have no essential link with emotion or any affective phenomenology. Thus, intuitions, as appearances, don't simply have representational content; they also *present* their content (Bealer 1998; Huemer 2005; Chudnoff 2013). Accordingly, we may again imagine Dud enjoying purely intellectual seemings in each of the evaluative Mary cases. As we saw in 3.1, it is clear there is still something he is missing, some way in which you more deeply appreciate the relevant values.

When the perceptual theorist tries to explain the epistemic benefit of emotion in terms of its status as evaluative appearance, she offers no way of saying what the *distinctive* epistemic contribution of emotion is. She offers no way of saying why it is epistemically better to have emotions rather than cold intuitions. Graham Oddie, for example, is open to this charge:

> An experience of the goodness of P … would be the state of P's *seeming* (*appearing*, *presenting itself as*) *good*, where this seeming is an experiential, non-doxastic take on the value of P. If there is such a state as an experience of the goodness of P, then, by analogy with the perceptual case, it would give me a reason to believe that P is good (2005, p. 40).

Oddie settles on desire as the relevant kind of evaluative experience, where he has in mind a kind of rich affective experience, rather than motivational states as such (personal communication). But if presentation is the justification-making feature of evaluative experience, then it doesn't matter what the evaluative experience *is*, not from the point of view of justification.

I conclude, then, that current articulations of the perceptual theory fail to provide resources for characterizing the distinctive epistemic significance of emotional experience.

2.4 UNDERSTANDING

There has been recent interest in the idea that emotions provide evaluative *understanding* (Roberts 2013, p. 46; Deonna and Teroni 2012, Ch. 10; Brady 2014, Ch. 4). Here, understanding p is not the same as knowing p or having a justified belief that p. I know

that Rome collapsed in the 5th century, and so do historians of Rome. But they understand this much better, and therefore their overall epistemic relation to the proposition 'Rome collapsed in the 5th century' is better than mine. Indeed, there is something initially plausible about the idea that the notion of understanding can explain what emotional experience provides. After all, it was natural to say that, compared to Dud, it's you who truly *grasps* the relevant values, that it's you who *really gets it*. I'll consider two versions of this view, and argue that both are lacking.

2.4.1 BRADY'S ACCOUNT

The view that emotional experience provides evaluative understanding is most thoroughly developed in Brady (2014). According to Brady, emotions promote evaluative understanding "through the capture and consumption of attention" (p. 118). The crucial epistemic feature of emotions is that they "can motivate us to search for reasons or evidence that bear on the accuracy of our initial emotional responses" (p. 129). To use Brady's example, imagine you awake to a noise in the middle of the night and feel afraid. Your fear motivates you to determine whether there is any genuine danger, and thus to determine whether you have reason to be afraid and to judge that you are in danger. Importantly, on Brady's view, when emotions motivate us to search for reasons, this often yields evaluative understanding.

What does Brady mean by 'understanding'? For starters, Brady writes that "achieving an understanding of emotional objects and events involves more than knowing that the objects and events have the evaluative properties that they do" (p. 139). Rather, understanding involves "a grasp or awareness of connections or links between

various items; it involves seeing how things fit together, how features are related, how facts support and explain other facts" (pp. 139-140).

Here, Brady follows the dominant view in the recent literature on understanding. For instance, consider Kvanvig:

> Understanding requires the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information. One can know many unrelated pieces of information, but understanding is achieved only when informational items are pieced together by the subject in question. … [U]nderstanding requires, but knowledge does not, an internal grasping or appreciation of how the various elements in a body of information are related to each other in terms of explanatory, logical, probabilistic, and other kinds of relations… (Kvanvig 2003, p. 192).

Thus, the idea here is that emotions motivate us to search for reasons, and this search typically reveals not merely *that* we are in danger, but *why*. When you hear the noise and fearfully search for signs of danger, you see the masked face in the window and thereby ascertain not only the danger but the danger-making features of your situation. And it is your fearful experience that facilitates this process.

Now, I grant that this is one epistemic benefit that emotions provide. But this view cannot fully explain the epistemic significance of emotional experience. The problem here is that Brady's view gives emotions a merely instrumental role, one that makes them dispensable. On this view, emotions are epistemically beneficial because of the *results* they *tend* to produce. Those results, however, can be achieved without emotions. Imagine that you are already awake and you see the masked face in the window. Imagine that you feel no fear. You merely intellectually recognize that you are in danger, and since you see the masked face and know what that means, you understand *why* you are in danger. It is still the case that when you become afraid, this makes the

dangerousness of your situation vivid to you. In your fear you somehow more fully register or appreciate the dangerousness.

Consider the evaluative Mary cases in light of this. We may suppose that Dud already fully understands *what makes* Genie's childhood a bad one. He can have this understanding as the result of careful reflection, without being motivated by his emotion to search for the relevant reasons. His careful reflection could have been motivated by any number of things—money, prestige, the love of wisdom. Regardless of what motivated him, we may suppose he does possess the relevant understanding. But this kind of understanding doesn't seem to fix Dud's limitation:

"Dud, you agreed Genie's childhood was terrible."

"Of course."

"But do you truly *understand* this?"

"Of course. It was terrible because of the isolation, the neglect, and so on."

Even though he understands *why* Genie's childhood was bad, you more fully register the badness in light of your being *moved* by the details of her childhood.

2.4.2 ROBERTS' ACCOUNT

I conclude that the instrumental role of emotion in acquiring understanding cannot fully explain the epistemic significance of emotional experience. This leaves open, of course, the view that emotional experience *directly* confers evaluative understanding, and not in virtue of motivating the search for reasons (Roberts 2013, Ch. 3). According to Roberts, sense perception often provides a similar benefit:

I can be told that the double-crested cormorant has a longer gular area than the neotropic cormorant, and, once I have learned what the gular area is, I have the wherewithal of a judgment. But I will certainly *understand* this judgment better if given the opportunity to see examples of the two species side-by-side (2013, p. 39).

And just as perception can deepen one's understanding of physical facts in this way, emotional experience deepens one's understanding of value (p. 52).

However, this view faces two challenges. The first is that in current epistemology, understanding it is always taken to be *holistic*. As in the Kvanvig passage from earlier, understanding requires grasping explanatory relations within a body of information (see also Grimm 2006 and Zagzebski 2001). But the sort of understanding Roberts describes here does not seem to be holistic in this way. Thus, an account is needed as to what kind of understanding this is, and why it is not holistic as other forms of understanding are taken to be.

The second problem has to do with *why* it is that emotional experience provides understanding. To see this, consider that Roberts thinks it is in virtue of their perceptual character that emotional experiences provide understanding, and he emphasizes their status as *appearances*. However, the same problem arises here as did for the perceptual theorist who emphasizes justification. After all, Roberts holds that the "emotionless person" does not enjoy the same evaluative understanding. However, we can imagine the emotionless person having intuitive seemings, where these are a kind of intellectual appearance. So, saying that emotional experience is a form of evaluative appearance can't explain why it provides understanding that is not otherwise available. Instead, it must be something about the way in which emotional experience is a form of evaluative

appearance. Plausibly, it would be something about the distinctive affective character of emotional experience, its thrills and chills. A full account should explain what exactly.

I conclude, therefore, that current accounts of emotions as sources of understanding cannot fully explain the epistemic significance of emotional experience.

2.5 CONCLUSION

In sum, I have argued that emotional experience generates some kind of epistemic good. This is on display in evaluative Mary cases. However, while numerous theories have been recently advanced, none of them are capable of fully explaining the epistemic good that you have but Dud lacks. I conclude, therefore, that there is a crucial aspect of emotional experience that remains unexplained. Of course, I have not addressed every possible or even every extant view, but we have seen that this explanatory failure bears out in examining several of the most influential recent accounts.

3.0 EMOTIONS AS REPRESENTATIONS OF VALUE

When you see a basket of puppies, you feel that familiar "melting" adoration. When you witness a police officer accepting a bribe, you feel shocked and appalled. But suppose things were switched around. You see a basket of puppies and feel shocked and appalled. You see a police officer accepting a bribe, and feel that melting adoration. Something is wrong with emotions that are switched around in this way, but what is it?

These examples make clear that at least some emotions can be assessed as correct, appropriate, or fitting. Feeling shocked and appalled is not a fitting response towards a basket of puppies. But what exactly is it for an emotion to be fitting or unfitting? Recently, many have answered as follows: When you are appalled by something, your emotional experience *represents* that thing as bad in some way. So, when you are appalled by a basket of puppies, your emotional experience is representing something as bad when that thing is in fact good. Conversely, when you experience adoration for something, you *represent* it as good in some way. So, when you adore police corruption, your emotional experience represents something as good when that thing is in fact bad. More generally, the idea is that emotional experience has evaluative representational content. Emotional experiences attribute evaluative properties to their objects. Accordingly, emotional experiences can be correct or incorrect, *in the same way* that beliefs, perceptual experiences, and indicative sentences can be correct or incorrect. Call

this **the Evaluative Content Thesis**. Here, the content of a state is whatever that state represents. A state has *evaluative* content *iff* that state represents an evaluative property such as goodness or badness. The goal of this chapter is to offer a novel argument in favor of the Evaluative Content Thesis.

The Evaluative Content Thesis has gained adherents in recent decades. In the theory of emotion, many have claimed that emotions have evaluative content, although there is dispute about whether emotions are beliefs (Foot 1978a, 1978b; Green 1992; Kenny 1963; Nussbaum 2001; Solomon 1976), thoughts (Greenspan 1988; Neu 1977), perceptual states (Döring 2003, 2009; Prinz 2004; Roberts 1988, 2003; Tappolet 2012; Johnston 2002), or sui generis mental states (Montague 2009).[21] In the philosophy of mind, representationalists should find the Evaluative Content Thesis attractive as part of their program for explaining phenomenal experience in terms of content (Byrne 2001; Dretske 1981, 1995; Hill 2009; Lycan 2003; Mendelovici 2014; Tye 1995, 2008, 2009). In metaethics, there is growing interest in the idea that, given that emotional experience has evaluative content, it can play a leading role in the epistemology of value, serving as a source of foundational justification for evaluative beliefs on the model of sense perception (Cuneo 2006; Dancy 2014; McDowell 1979, 1985; Pelser 2014; Tolhurst 1990).

However, in spite of its broad appeal, the Evaluative Content Thesis has been the object of recent criticism (Brady 2014; Deonna and Teroni 2012; Schroeter et al. 2015), and three alternative views are worth noting here. First, the most austere alternative to the Evaluative Content Thesis is the claim that emotional experience has no representational

---

[21] A related view is the claim that *desires* have evaluative content (Oddie 2005; Stampe 1986, 1987; Schafer 2013). I will mostly set this view aside in this chapter, except for in section 3.1.1.

content whatsoever. This view, which we may call **the No Content Thesis**, can reasonably be attributed to Hume.[22] Second, a more modest view is that emotional experience has some content but that its content is *merely* about one's bodily states (such as a racing heart or a surge of adrenaline). This view, which we may call **the Body Content Thesis**, was famously held by William James.[23] Finally, a third alternative to the Evaluative Content Thesis we may call **the Propositional Content Thesis**. Notice that emotions often take propositions as their objects. One can be happy, sad, angry, or amused *that p*. According to the Propositional Content Thesis, whatever representational content an emotion has is exhausted by the proposition that is the object of the emotion.[24] This view has recently been defended in Deonna and Teroni (2012; 2015; forthcoming).

This chapter offers challenge to both proponents and critics of the Evaluative Content Thesis. As concerns its proponents, I will show that an influential argument for the Evaluative Content Thesis is not in fact persuasive. Many proponents of the Evaluative Content Thesis have held that the fittingness of emotion directly favors their view. I'll show that this is not correct. On the other hand, against the critics of the Evaluative Content Thesis, I will argue that there is a way in which the fittingness of emotion does figure into an argument for the Evaluative Content Thesis, a novel argument that it is the goal of this chapter to develop. By way of preview, my argument will run as follows:

---

[22] "A passion is an original existence, … and contains not any representative quality, which renders it a copy of any other existence or modification. When I am angry, I am actually possest with the passion, and in that emotion have no more a reference to any other object, than when I am thirsty, or sick, or more than five foot high" (*Treatise* 415). The details of Hume's actual view are, of course, controversial.
[23] James (1884); Damasio (1994).
[24] This view will need to say something about emotions that are on the face of it *not* about propositions, such as one's fear of *the growling dog*.

1. Emotions can be individuated by their fittingness conditions. For instance, fear is fitting *iff* it is directed at potential harms or threats, while grief is fitting *iff* directed at losses.

2. Emotions can be individuated by their phenomenal character. For instance, when you experience fear, you are not likely to mistakenly judge that you are *amused*. That's partly because there are signature differences in how fear and amusement *feel*.

3. The fact that both fittingness conditions *and* phenomenal character can differentiate emotions is an implausible coincidence, unless there is some explanatory link between fittingness conditions and phenomenal character.

4. The best accounts of this explanatory link commit us to the Evaluative Content Thesis.

5. So, we should accept the Evaluative Content Thesis.

Let's call this **the Argument from Massive Coincidence**. What we may call **the Massive Coincidence** is the fact that, as fittingness conditions vary from one emotion type to another, so does phenomenal character; as phenomenal character varies, so do fittingness conditions. I'll show that the neither the No Content Thesis, the Body Content Thesis, nor the Propositional Content Thesis can successfully explain the Massive Coincidence.

The rest of this chapter is divided up into four sections, each of which defends one of the four premises of the Argument from Massive Coincidence.

3.1 FIRST PREMISE: THE FITTINGNESS OF EMOTION

In this section, I'll show that the fittingness of emotion does not directly establish the Evaluative Content Thesis, as many authors have held. After that, I will begin the Argument from Massive Coincidence by arguing for premise 1.

As we saw at the outset, emotions can be assessed for fittingness. Furthermore, we may now observe that the fittingness of an emotion in some way depends on whether the object of the emotion instantiates the relevant evaluative property. For instance, anger is fitting when directed at the offensive, fear when directed at the dangerous, and admiration when directed at the admirable. Being repulsed by a cute baby, or afraid of a pencil, or proud of one's evil deeds, would be unfitting (barring special contrived circumstances), since these objects do not instantiate the relevant evaluative properties.[25]

By *fittingness*, I mean a generic notion of appropriateness or correctness. There is a debate about whether emotions have their own proprietary species of fittingness (D'Arms and Jacobson 2000), or whether the fittingness of emotion is always *moral* fittingness (Roberts 2003, according to D'Arms and Jacobson). For instance, for D'Arms and Jacobson, amusement at an offensive but funny joke is fitting even though morally inappropriate, and perhaps inappropriate *overall*. In this chapter, I mean to use a notion of

---

[25] I grant that sometimes the fittingness of an emotion also depends on non-evaluative facts. If you are angry that John deleted your manuscript, this is fitting only if John actually did this. Whether this is generally a requirement on the fittingness of emotion is complicated by the fact that we can have fitting emotions towards non-real objects, such as fictions and the contents of imaginative states.

fittingness that is neutral on that issue. I mean to invoke a generic notion of correctness. My starting point is simply the observation that to be appalled by a cute puppy is inappropriate or incorrect in *some* way. Any more specific notion of emotional fittingness has to be argued for.

Many authors seem to hold there is a simple, direct inference from the fittingness of emotion to the Evaluative Content Thesis. If that's so, then the Argument from Massive Coincidence is completely superfluous. Thus, before arguing for premise 1, I want to argue that there is no *direct* inference. Rather, the fittingness of emotion can only form part of an overall case for the Evaluative Content Thesis.

3.1.1 THE DIRECT INFERENCE: WHY IT FAILS

Many hold that the fittingness of emotion *directly* supports the Evaluative Content Thesis:

> Fear is a response to one's situation being of a certain kind: something is dangerous. As such, the rational appropriateness of fear depends on the situation's really being that way: fear is in part a cognitive state with mind-to-world direction of fit. Yet fear is also evaluative…. The content of fear is not merely that something is about to be destroyed or harmed, but that this is a bad thing (Helm 2001, p. 6).

> [Fear's] negative evaluative content is needed to explain … why fear amounts to a *reasonable* reaction in certain situations (Greenspan 1988, p. 3).

> We are prone to assess our emotions with respect to how they fit evaluative facts. We criticize our fears when they are about things that are not fearsome, for instance. This practice suggests that the object of fear is represented as fearsome (Tappolet 2012, p. 210).

However, the mere fact that emotions can be fitting or unfitting does not show they have evaluative content. *Actions* can be fitting or unfitting, appropriate or inappropriate, etc., but actions do not have representational content. If actions had representational content, then they could be assessed in alethic terms such as truth, accuracy, veridicality, etc. But actions cannot be assessed in alethic terms (except perhaps in special cases where actions – a code of gestures, for instance – have been artificially imbued with semantic value). For instance, if your house is on fire, it is fitting for you to rush out with your most treasured possessions. But rushing out with your most treasured possessions is not *true* (or accurate, veridical, etc.). So, the direct inference from fittingness to representational content fails, since actions can be fitting but do not have representational content.

One could defend the direct inference by claiming that actions actually *do* have content: actions are partly constituted by intentions, and intentions have content. Suppose you are standing in the street and waving. Whether you are hailing a taxi or saying hello depends on your intention. The bodily motion is the same. Thus, actions have content since they are constituted by a state with content.

However, the only non-controversial sense in which intentions have content is the following: when you intend to Φ, Φ is in some way the content of your intention. But this is not *evaluative* content. Now, one might appeal here to the so-called "guise of the good" thesis, the view that the desire to Φ represents Φ-ing as *good* (e.g., Stampe 1987). Using this thesis, one might reason as follows: Since desires partly constitute intentions, the intention to Φ represents Φ-ing as good; since intentions partly constitute actions, the act of Φ-ing also represents Φ-ing as good. Accordingly, one might hold that, when

actions are unfitting, it is because they have a false evaluative content. This defends the direct inference by showing that actions are not in fact a counterexample.

However, what we have here is, if anything, a problem for the guise of the good thesis. If the guise of the good thesis ends up committing us to the claim that actions can be true or false, accurate or inaccurate, etc., then that is a problem for the guise of the good thesis, since we simply do not assess actions in any such alethic terms. Proponents of the guise of the good thesis need to resist the inference from the evaluative content of desire to the evaluative content of action. Indeed, such an inference does seem mistaken. The mistake lies in the claim that, because actions are partly constituted by something with content, actions must also have content. We should reject this. If I upholster my chair with cosmological maps from the middle ages, this does not mean that my chair is inaccurate, even though it is partly constituted by something that is.

In sum, since actions cannot be assessed in alethic terms, actions do not have representational content. So, actions have fittingness conditions without being representational. This shows that the mere fact that emotions have fittingness conditions does not entail that emotions have representational content. The direct inference from fittingness to content fails.

3.1.2 FITTINGNESS CONDITIONS INDIVIDUATE EMOTIONS

The discussion above shows the fact that emotions are assessable for fittingness does not *directly* support the Evaluative Content Thesis. However, the fittingness of emotion can form *part* of a case for the Evaluative Content Thesis. Here, I will argue for premise 1 of the Argument from Massive Coincidence: Emotions can be type-individuated in terms of

their fittingness conditions. This premise can be seen to be plausible just by considering examples. Fear is fitting when directed at the dangerous, disgust when directed at the foul, pride when directed at one's achievements, envy when directed at the enviable, admiration when directed at the admirable, amusement when directed at the comical, and so on. Each emotion has a unique set of fittingness conditions that differentiates it from other emotions.

Notice that neither does *this* directly support the Evaluative Content Thesis. For, emotions can be differentiated in terms of fittingness conditions, without saying that emotions *represent* their fittingness conditions. So, we should reject reasoning like the following:

> [Fear's] negative evaluative content is needed to explain why it amounts to fear rather than some other reaction to an envisioned possibility, such as thrilled anticipation (Greenspan 1988, p. 3).

It simply isn't true that evaluative content is needed to differentiate fear from other emotions. Fittingness conditions alone can do this.

In sum, emotions can be type-individuated in virtue of their fittingness conditions. In the next section, I will show that there is a second way to type-individuate emotions, namely, in virtue of their phenomenal properties.

3.2 SECOND PREMISE: THE PHENOMENOLOGY OF EMOTION

William James writes that "the internal shadings of emotional feeling ... merge endlessly into each other. Language has discriminated some of them, as hatred, antipathy, animosity, dislike, aversion, malice, spite, vengefulness, abhorrence, etc., etc."; but, James adds, those "internal shadings" are primarily distinguished from each other "by their conscious subjective tone" (James 1890, p. 448). Here, James is claiming that emotions can be type-individuated based on what it is like to experience them. Emotions have, we might say, *phenomenal signatures* that tell them apart. The strongest version of this view – and the most frequently discussed – quantifies universally:

> *The Uniqueness Thesis*: All emotions can be type-individuated by their phenomenal properties.

The Uniqueness Thesis has a venerable history of detractors (e.g., Lyons 1980; Nussbaum 2001; Ryle 1949; and Schachter and Singer 1962). However, in this section, I will show that the Uniqueness Thesis (which corresponds to premise 2 in the Argument from Massive Coincidence) is nevertheless plausible. First, I will offer two positive arguments in favor of the Uniqueness Thesis – an enumerative induction, and an argument from self-knowledge in the absence of other evidence. Second, I'll address several influential objections against the Uniqueness Thesis.

3.2.1 THE ENUMERATIVE INDUCTION

Everyone should accept that there are *some* emotions that can be distinguished from each other on the basis of their characteristic feelings. For instance, if you are *enraged* by something, you are not likely to mistakenly judge that you are experiencing *amusement*. These emotions have obvious differences in their phenomenal character. In general, we can distinguish between emotions that feel pleasant and those that feel unpleasant. In addition to pleasantness, emotions can also be assessed for felt *intensity*. And some emotions can simply be distinguished from each other on this basis. Rage, for instance, feels more intense than indignation. And if we consider in tandem the pleasantness and intensity of feelings, we can distinguish many more emotions from each other. For example:

Elation – pleasant, high intensity.

Contentment – pleasant, low intensity

Terror – unpleasant, high intensity.

Boredom – unpleasant, low intensity.

Every emotional experience, it seems, can be located along these two dimensions (Russell 1980). However, pleasantness and intensity do not exhaust emotion phenomenology. Terror, rage, and despair are unpleasant, and each can be comparably intense, but still there are obvious *qualitative* differences between these states. Despair has a sinking, crushing feel. Terror has an intense, fluttery, focused feel. Rage feels hot

and dense. Of course, these descriptions inevitably fall short, as will any attempt to verbally characterize emotional phenomenology. But that is simply due to the limitations of language. In verbally characterizing emotional phenomenology, the best we can do is *point* to the feelings in question.

These considerations show that many emotions can be distinguished on the basis of what they feel like. Accordingly, a straightforward enumerative induction supports the Uniqueness Thesis. Of course, such an induction is not wholly conclusive. There *could* nevertheless be an emotion without a phenomenal signature. But exceptions are possible on any enumerative induction.

I grant, moreover, that some emotions do not admit of obvious phenomenal differences. For instance, consider shame and embarrassment, awe and wonder, and gratitude and admiration. These emotion pairs do not have obvious phenomenal differences. However, the Uniqueness Thesis does not claim that all emotions have *obvious* phenomenal differences. The phenomenal differences between the emotions in the above pairs are subtle, and therefore these emotion pairs are ill-suited as the *basis* of an inductive case. Nevertheless, their subtlety does not undermine the inductive case.

3.2.2 THE ARGUMENT FROM SELF-KNOWLEDGE

The enumerative induction can be supplemented by a further inductive consideration, namely, the argument from self-knowledge. When we experience an emotion, we often know what emotion we are experiencing. When I reach the mountaintop and I am moved by what I see, it is evident to me that I am experiencing longing rather than, say, grief or amusement. How *do* know that I am experiencing longing rather than some other emotion? It is natural to answer that I know this based on my feelings. What I am

experiencing manifestly feels like longing. It doesn't feel like grief, amusement, or some other emotion.

If this is correct, it would show that my emotion has a proprietary feeling, not shared by other emotions. Otherwise, the feeling of the emotion would not be sufficient evidence for determining that I am experiencing longing rather than some other emotion. So, if emotional self-knowledge is sometimes based on feelings alone, then this shows that in such cases the feeling of the emotion is type-individuating.

An alternative explanation of emotional self-knowledge runs as follows. One's feelings play only a limited evidential role. In the mountaintop case, perhaps the pleasant valence of one's feelings alerts one to the fact that some positive emotion is underway. But in order to determine what specific emotion one is having, one must draw on further evidence. Gilbert Ryle was perhaps the most influential proponent of this view. As Ryle put it, "[p]ains do not arrive already hallmarked 'rheumatic', nor do throbs arrive already hallmarked 'compassionate'." And since the feeling of an emotion doesn't "arrive hallmarked", one must use other evidence to self-ascribe the emotion. Thus, "[w]hen a person reports a chill of disquiet or a tug of commiseration, he is not merely reporting a feeling; he is giving a diagnosis of it….". And as each emotion shares its phenomenology with other states, it is easy to "diagnose as a twinge of remorse what is actually a twinge of fear", or to "ascribe to dyspepsia a feeling which is really a sign of anxiety" (Ryle 1949, p. 90).[26]

_____

[26]Ryle's view anticipates the famous study in Schachter and Singer (1962). In Section 3.2.3, I'll discuss why I find no challenge in the results of that study.

In the mountaintop case, what further evidence might one be drawing on to "diagnose" one's feelings as a state of longing rather than some other state? Two salient factors are, first, one's environment, and second, one's behavior. In the mountaintop case, one is in an environment in which longing is typical, even paradigmatic. Moreover, one can observe one's own behavioral cues – one's attention is fixed, one wants to stay put, to snap a photograph, and so on.

Accordingly, a case that would decide between these two explanations would be a case in which the subject has a conscious emotional experience in addition to emotional self-knowledge, but lacks access to these further sources of evidence. Fortunately, such cases are available. One example comes from Michael Gazzaniga's classic research on split-brain patients. One patient, V.P., was shown a video of violent acts, such as someone shoving a stranger off a balcony, or throwing a Molotov cocktail. But this footage was only shown to the right hemisphere, where V.P.'s linguistic abilities were almost entirely deficient. However, V.P.'s right hemisphere was still communicating with his linguistically competent left hemisphere:

> V.P.: I don't really know what I saw. I think I just saw a white flash.
>
> Experimenter: "Were there people in it?"
>
> V.P.: "I don't think so. Maybe just some trees, red trees, like in the fall."
>
> Experimenter: "Did it make you feel any emotion?"
>
> V.P.: "I don't really know why, but I'm kind of scared. I feel jumpy. I think maybe I don't like this room, or maybe it's you, you're getting me nervous."
>
> V.P. turned to female experimenter and in private said: "I know I like Dr. Gazzaniga but right now I'm scared of him for some reason" (Gazzaniga and Smiley 1984, p. 203).

V.P. identified his emotion as fear or nervousness, it seems, simply on the basis of what he was feeling, even though he had no idea what he was afraid of. V. P. was not in a scenario typical of fear. He can be seen casting about to find a suitable object of fear to explain why he feels as he does. This suggests that he identifies his emotion as fear prior to ascertaining what his fear is about and whether fear is a typical response to his environment. Moreover, V. P.'s behavior is not especially salient in this case. He is neither fighting nor fleeing. The primary behavior on display is that of searching for an explanation of one's emotional state. But this is not a behavior distinctive of fear. All V.P. has to go on, then, is what he feels.[27]

I conclude, therefore, that the feeling of an emotion is often sufficient evidence for knowing what emotion one is having. This provides inductive evidence, which, in addition to the enumerative induction, shows that the Uniqueness Thesis is highly plausible. Why, then, have so many rejected the Uniqueness Thesis? In the next section I'll address several of the most influential objections.

---

[27] I take it that what I say here is compatible with a range of theories about the nature of self-knowledge. Even Carruthers (2011), who channels important aspects of Ryle's view in holding that self-knowledge of propositional attitudes is always indirect or inferential, allows that phenomenal features of affective states can form the basis of emotional self-knowledge (see Carruthers 2011, Ch. 5). For Carruthers, the knowledge that one is afraid can be based on feelings, but the knowledge that one is afraid *that p* cannot (not fully, at any rate). Even Gallois (1996), who holds that self-knowledge of mental states is in some sense based on the standards of rationality for those states, allows that emotions can be individuated by their respective feelings, and that accordingly such feelings can sometimes serve as the basis of emotional self-knowledge. These two accounts are especially worth mentioning, since, out of all the accounts of self-knowledge, these two seem to be the least likely to end up being friendly to the claim that feelings can serve as the basis of emotional self-knowledge.

3.2.3 OBJECTIONS TO THE UNIQUENESS THESIS

In this section, I will address five influential objections against the Uniqueness Thesis, and show that their influence is underserved.

First, one influential objection comes from the famous study in Schachter and Singer (1962). Allegedly, subjects induced with the same feelings self-ascribed different emotions based on environmental differences. However, this study failed to replicate on numerous occasions (e.g., Marshall 1976, Marshall and Zimbardo 1979, and Maslach 1979).[28] The study is therefore widely discredited and poses no empirical challenge to the Uniqueness Thesis.

A second, perhaps tempting objection is that one can have an emotion without experiencing any feelings at all, as with non-conscious or long-term emotions. However, this objection has no traction with the Uniqueness Thesis. The Uniqueness Thesis doesn't say that emotions are always conscious, but that *when* they are conscious, they have distinctive feelings. Instead, what this objection is relevant to is the feeling theory of emotion, which *identifies* emotions with their respective feelings.[29] The feeling theory entails the Uniqueness Thesis, but not vice versa. In this chapter I remain neutral on the feeling theory.

A third objection is that, if feelings provide sufficient evidence for self-ascriptions of emotion, then there is a puzzle as to how it is that people are sometimes mistaken about what emotion they are feeling. However, this is not difficult to explain, although we should not expect just one answer. In some cases, people have motivations for

---

[28] For more thorough discussion, see Prinz (2004, pp. 70-71). Failure to replicate aside, the Schachter and Singer study has also met with several devastating objections, notably in Griffiths (1994, p. 82-83), Levenson (2003, pp. 212-213), and Reisenzein (1983).

[29] For further criticism of the feeling theory, see Sizer (2006), which develops a novel line of objection.

actively ignoring the evidence of their feelings, as in cases of self-deception. In other cases, people are distracted or imperfectly attentive to their feelings. In still other cases, emotions have *subtle* differences in their phenomenal signatures, as in the case of wonder and awe. Awe involves a shrinking paralyzed feeling when compared with wonder, which is more like an outward-reaching fascination (Nussbaum 2001, p. 54, n. 53). But this phenomenal difference is subtle. So, people might easily be mistaken in such a case, if they are only judging on the basis of what they feel. Moreover, some individuals are just not skilled at classifying what emotion they are experiencing. The mere fact that two states are phenomenally different does not mean we will always be good at detecting and describing the phenomenal difference.

A fourth objection is that emotions are distinguished on the basis of their *logic*, rather than their phenomenal properties. For instance, consider shame and embarrassment. Robert Solomon writes that "the feelings and sensations involved in the two emotions are of little relevance in discriminating between [them]". Instead, we discriminate between shame and embarrassment based on "the 'logic' of the situation" (Solomon 1976, p. 98). In embarrassment, we are put in an awkward position but are not responsible for it, while in shame we are responsible for whatever the evil is that has arisen. However, it is compatible with the Uniqueness Thesis that there are other ways to individuate emotions besides their respective feelings. Moreover, there *are* phenomenal differences between shame and embarrassment, even though they are subtle. For starters, shame is typically more painful than embarrassment, and feels *heavier*. Embarrassment flushes the face while shame sinks in the guts. In embarrassment one feels stiff and squirmy; in shame one feels cowered down.

Fifth, and finally, it is often objected that some emotions share their bodily phenomenology with each other. For instance, Solomon writes that "feelings are never sufficient to identify or to differentiate emotions…", since we can have "all the symptoms of emotionality, for example, flushing, pulsing" without having any emotion (Solomon 1976, p. 99; see also Nussbaum 2001, pp. 60-61). However, while emotional experience often co-occurs with bodily feelings, bodily feelings are not the only phenomenal properties the Uniqueness Thesis can appeal to. Here, I understand bodily feelings in terms of *felt location*: A feeling is bodily *iff* the feeling is in some way experienced as *located* in the body. So, a stomachache is felt to be located in the stomach. Even a fever, which doesn't have a specific location, is felt to pervade the whole body. Following Stocker (1983), I'll call non-bodily feelings *psychic*. So, psychic feelings are feelings that are not bodily, which means they are not experienced as being located in the body.

While I can't offer a systematic argument here, it is plausible that there are psychic feelings. For starters, there are a variety of psychological disorders that are best explained in terms of phenomenological disturbances that require some notion of psychic feelings. For instance, in depersonalization, many subjects report that the world as they visually experience it is unreal or dreamlike.[30] These reports are extremely robust, and can't be explained away as mere hyperbole (Guralnik et al. 2000; Hunter et al. 2004; Simeon et al. 1997; Simeon et al. 2000). Are these subjects simply reporting a *judgment*? That is implausible, since the subjects have 'intact reality testing', meaning they know

---

[30] These symptoms are called "derealization" (see Sierra et al. 2002). To date, there is no published, properly-thorough catalogue of derealized patients' reports of their symptoms, but for some examples refer to the Cambridge Depersonalization Scale in Simeon and Abugel (2006), pp. 75-76. Helpful overviews of depersonalization can be found in Sierra (2009), Simeon (2004), and Simeon and Abugel (2006). Note that not all subjects with depersonalization report derealization symptoms.

they are experiencing a disturbed form of consciousness. They know they are not actually in an unreal dream-world. So, depersonalized subjects are reporting an *experience*. Plausibly, then, these subjects have either lost a feeling of reality that normal observers have, or they have gained a feeling of unreality. But there is no evidence suggesting these feelings of unreality are experienced as located in the body. For instance, there is no evidence that these subjects have abnormalities in the autonomic nervous system (which controls a range of bodily processes such as arousal). Indeed, it is difficult to see how such feelings could be specific to a bodily region.[31]

Given that there are psychic feelings, a proponent of the Uniqueness Thesis can take the view that emotional phenomenology is at least partly constituted by psychic feelings. Thus, even if, as Solomon claims, some emotions share their *bodily* feelings, this wouldn't rule out that those emotions can be individuated by their respective *psychic* feelings.

In sum, I've offered two arguments for the claim that emotions have type-individuating feelings, an enumerative induction and an inductive argument from self-knowledge. I have also argued that the most influential objections against the Uniqueness Thesis can be overcome. Thus, it is plausible that emotions can be individuated by their phenomenal properties. And as we saw in the previous section, emotions can also be individuated based on their fittingness conditions. With these two claims in place, we now have the materials needed to press the Argument from Massive Coincidence. Why, after all, can emotions be type-individuated in these two ways? Is it a mere coincidence,

---

[31] For further examples of psychiatric disorders that warrant the notion of psychic feelings, see Bayne (2011), McLaughlin (2009), and Ratcliffe (2008).

or is there some explanatory relation between an emotion's fittingness conditions and its phenomenal character? In the next section, I will argue the latter.

3.3 THIRD PREMISE: IS THE MASSIVE COINCIDENCE A MERE COINCIDENCE?

In this section, I'll argue for premise 3. Could it just be a coincidence that both fittingness and phenomenal character differentiate emotions? Certainly, coincidences happen. But it is a stunning coincidence that all emotions can be differentiated by fittingness conditions *and* phenomenal character.

Indeed, it is plausible that the Massive Coincidence is a *necessary* fact, and the fittingness conditions and the phenomenal character of an emotion co-vary across possible worlds. There is no possible world in which anger is a fitting response to, say, goods contained in one's future (as such). There is no possible world in which hope is a fitting response to, say, the comical (as such). There is no possible world in which fear is the fitting response to, say, the achievements of one's children (as such). Of course, there are possible worlds in which the achievements of one's children *happen to indicate* danger. There are worlds in which children use their achievements to *harm* their parents. But that's what the "as such" is for. Fear can be a fitting response to the achievements of one's children *as likely dangers* but not *as* the achievements of one's children. So, the fittingness conditions of an emotion are a necessary feature of it. Plausibly, an emotion's phenomenal character is also a necessary feature of it. There is no possible world in which anger feels like amusement, or relief feels like terror, or grief like gratitude. As Tim Crane puts it, we should not accept "the possibility that there is a world in which

contentment feels to someone as anxiety feels to me" (Crane 1998, p. 9). So, not only do emotions have type-individuating phenomenal signatures, but they have them necessarily. So, *in every possible world* emotions can be differentiated both by fittingness conditions and by phenomenal character. We should expect an explanation of this.

One might object here that necessary facts do not need to be explained. They are necessary! Or rather, if some fact obtains in the actual world, and we discover that it is necessary, that *is* the explanation of why it obtains. There is no further question as to what explains the necessity. So, the necessity of the Massive Coincidence removes, rather than intensifies, the need for an explanation.

However, it is just not correct that necessities never need to be explained. The proposition (A) '2 + 2 = 4' is a necessary fact. So is the proposition (B) 'either 2 + 2 = 4, or pigs can fly'. But it is plausible that A in some sense explains B. Moreover, it is commonplace that there are explanations of mathematical facts.[32]

Perhaps the more pressing question is why we should expect these two necessary properties of emotion to figure in the *same* explanation, rather than being fully explained by independent, unrelated factors. Why should we expect a *unifying* explanation? In response, I grant that there *can* be non-unifying explanations. However, it is plausible that, all things being equal, a unifying explanation is preferable to a non-unifying explanation. I take this to be a fact about explanation in general. That is why, for instance, the development of a unifying explanation of electricity and magnetism marked a feat of explanatory progress. As Michael Friedman remarks:

---

[32] See the conversation with mathematician Michael Atiyah in Gray and Wilson (2001), especially pp. 18-20.

> [S]cience increases our understanding of the world by reducing the total number of independent phenomena that we have to accept as ultimate or given. A world with fewer independent phenomena is, other things equal, more comprehensible than one with more (Friedman 1974, p. 15).

To say this is not to endorse a unificationist theory of explanation (Friedman 1974, Kitcher 1981). It is simply to say that unification is one explanatory virtue among others. This allows that in deciding which explanation is best, we must consider other explanatory virtues as well, such as breadth or predictive power. But if two explanations are equal in these virtues, then a difference in unification can decide between them.

Moreover, the virtue of unification does not simply apply to explanations of contingent causal phenomena, since necessities are similarly in need of explanation, albeit, perhaps, explanations of a different sort. Presumably, a unifying explanation of the Massive Coincidence will involve some notion of ontological or metaphysical dependence, rather than causation. For instance—and this is just an example—one unifying explanation would have it that both the fittingness conditions of an emotion and its phenomenal character are ontologically dependent on its evaluative content. Now, how best to understand the notion of ontological dependence here is as important as it is complicated. Perhaps it is best understood in terms of *grounding* (Fine 2001). However, for our purposes such matters need not be decided. At present, the question is why we should expect a unifying explanation of the fittingness conditions and phenomenal character of an emotion, and my response is just that unification is in general an explanatory virtue. The question is, then, are there plausible unifying explanations of the Massive Coincidence? In the next section I will turn to this.

3.4 FOURTH PREMISE: EXPLAINING THE MASSIVE COINCIDENCE

There are three initially viable unifying explanations that are worth considering. The first two entail the Evaluative Content Thesis, while the third is totally independent of the Evaluative Content Thesis. I want to show that the unifying explanations that entail the Evaluative Content Thesis are better off than the unifying explanation that doesn't.

For starters, it is easy to see how, if emotions have evaluative content, this would explain their fittingness conditions. An emotion's fittingness conditions would simply *be* accuracy conditions: an emotion is fitting *iff* it correctly attributes an evaluative property to its object. It is in virtue of attributing an evaluative property that an emotion becomes assessable for correctness. As is clear, on the Evaluative Content Thesis, the explanatory relation between evaluative content and fittingness is simple and direct.

Importantly, in order for the representational content of emotion to determine its fittingness conditions, it would need to be *evaluative* content. Otherwise, it is difficult to see how an emotion's content could explain its fittingness conditions. If, say, the Body Content Thesis is correct, and the content of anger is merely that my blood pressure is rising, etc., then this content wouldn't explain why anger should be felt only in response to wrongs and offenses. Alternatively, if the Propositional Content Thesis is correct, this, too, would be unable to explain the fittingness conditions of emotion in terms of representational content. (It may be able to explain fittingness conditions in some other way, but not in terms of representational content.) After all, different instances of anger can have very different propositions as their objects, but the evaluative property in their fittingness conditions remains the same. One's anger that Suzy took the last beer and one's anger that Leo didn't win the Oscar are both fitting only if these events are wrongs

or offenses. Thus, the Body Content Thesis and the Propositional Content Thesis are not able to explain the fittingness of emotion in terms of representational content, because they do not appeal to *evaluative* representational content.

So far, we've seen that evaluative content can explain the fittingness of emotion. But in order to find a *unifying* explanation, there must also be an explanatory relation between evaluative content and the phenomenology of emotion. There are two possible views here, depending on whether evaluative content or phenomenology has explanatory priority.

On the first view, evaluative content explains emotional phenomenology. Most representationalists, say, about perceptual experience hold a view with this structure. Content determines both the phenomenal character of perceptual experience, and its correctness conditions (e.g., Siegel 2011). On this view, then, it is in virtue of representing dangers that fear feels as it does, and has the fittingness conditions it does. Of course, more would need to be said to articulate the way in which content explains phenomenology. But accounts of this have been thoroughly developed by representationalists with respect to other aspects of consciousness, such as perceptual experience. Proponents of the Evaluative Content Thesis have the advantage of being able to draw on such accounts in characterizing the way in which the evaluative content of an emotion explains its phenomenology.

An alternative view reverses the order of explanation, and holds that it is the phenomenal character of an emotion that explains its evaluative content. This account relies on the notion of *phenomenal intentionality* (Chalmers 2004, 2006; Kriegel 2007, 2013; Horgan and Tienson 2002; and Loar 2003). A state has phenomenal intentionality

iff that state has intentional properties as well as phenomenal properties, and that state has (at least some of its) intentional properties in virtue of (at least some of its) phenomenal properties. Those who claim that some states have phenomenal intentionality typically have the following intuition: If there were a brain in a vat who was my phenomenal duplicate, it would also *thereby* share many of my intentional properties as well. Thus, advocates of phenomenal intentionality propose a break with the dominant view that content is determined by external, environmental factors.[33]

Could emotional experience have phenomenal intentionality? Assuming that *some* states have phenomenal intentionality, it is plausible that emotions could as well. For instance, recall that emotional experience can be assessed in terms of valence and intensity. Perhaps the valence of an emotion determines whether it represents its object as good or bad, while the intensity of an emotion determines *how* good or bad the emotion represents its object as being. So, being slightly upset about F represents F as somewhat bad, but being enraged about F represents F as being very bad. That explains why, if F is only a minor infraction, rage is unfitting—even though rage correctly represents F as bad in some way, rage attributes *too much* badness.

But an emotion's fittingness doesn't simply depend on whether and to what degree the emotion's object is good or bad. Rather, emotions have more determinate fittingness conditions than that. Anger and disgust can be comparably intense, and both are negatively valenced, but they still have differing fittingness conditions. That's why you should be disgusted by cockroaches but not angry with them. So, if an emotion's fittingness conditions are determined by its content, then emotions must have highly

---

[33] Of course, advocates of phenomenal intentionality are not the only ones who believe in *narrow content*, content that isn't determined by external factors. For instance, some hold that functional properties can determine content independently of one's environment (e.g., Harman 1982; Block 1986).

determinate evaluative content, since an emotion's fittingness conditions are highly determinate. It can't be the case that emotions simply represent their objects as having some degree of goodness or badness.[34] Thus, if an emotion's phenomenal character determines its content, then it must be in virtue of suitably determinate features of its phenomenal character. Fortunately, as we have seen in Section 3.2, emotional phenomenology is highly determinate. Each emotion has a determinate phenomenal signature. So, on the present account, it must be in virtue of an emotion's determinate phenomenal "tone" that the emotion has the determinate evaluative content it does. So, when you are moved with longing by the view from a mountaintop, your emotional experience represents the scene as sublime. And it has this content in virtue of the distinctive qualitative character of felt longing. On this view, then, an emotion's phenomenal character determines its evaluative content, and its evaluative content determines its fittingness conditions.

We have seen two initially viable accounts that use the evaluative content of emotion to unify its fittingness conditions and phenomenal character. Are there unifying explanations that do not make use of evaluative content? Such an explanation has recently been developed in Deonna and Teroni (2012; 2014; 2015; forthcoming). Deonna and Teroni hold that emotions are "experiences of our body as ready or poised to act in various ways towards an object" (Deonna and Teroni 2012, p. 80). The claim here is that an emotion is constituted by the feeling of one's body preparing to respond to what the emotion is about. In this sense, emotions are "feelings of action readiness" (Frijda 1986). Accordingly, emotions can be type-individuated by the sorts of actions they prepare the

---

[34] This opposes Tye (2008), which claims that emotional experience represents less determinate evaluative properties like good and bad (although Tye allows that fear represents dangers).

body for. For instance, with fear, the associated action type is that which will "contribute to the neutralization of what provokes the fear" (Deonna and Teroni 2012, p. 80). While one may take any number of specific courses of action towards the growling dog, there is a behavioral theme, namely, that of neutralizing the dog's impact. Thus, Deonna and Teroni hold that the action types associated with emotions are (i) determinate enough to individuate emotions, yet (ii) open-ended enough to allow for the diversity of specific actions one might engage in while emotional.

Crucially, for Deonna and Teroni, *an emotion's proprietary action type is what determines its fittingness conditions*. They agree that, in fearing the growling dog, the fittingness of one's fear depends on the dangerousness of the dog. But this is so, they claim, "precisely because [fear] consists in feeling the body's readiness to act so as to diminish the dog's likely impact on it…, and this felt attitude is correct if and only if the dog is dangerous." Similarly, they agree that the fittingness of anger towards someone depends on her offensiveness. But this is so, they claim, "precisely because it consists in feeling the body's readiness to act so as to retaliate one way or another, and this felt attitude is correct if and only if the person is or has been offensive." These remarks are meant to explain "the fact that distinct emotions are correct only when their object exhibits a specific evaluative property" (Deonna and Teroni 2012, p. 81).

The idea here is that the fittingness conditions of the action type associated with an emotion explain the fittingness conditions of the emotion itself. Emotions, we might say, *inherit* their fittingness conditions from their associated action types. And the phenomenology of emotion just is the phenomenology of preparing to engage in the relevant action type. So, the phenomenology of an emotion explains its fittingness

conditions, without any reference to evaluative representation. This view we may call **the Motivation Thesis**.

However, there is good reason to reject the Motivation Thesis. It isn't clear that the action types associated with emotions are fine-grained enough to track the fittingness conditions associated with emotions. For instance, consider fear and disgust. Fear is fitting when directed at the dangerous, disgust when directed at the foul. But both of these emotions seem to motivate neutralizing the impact of the emotion object. Against this, Scarantino (2014, p. 181) claims that in fear, one wants to *avoid* the object, while in disgust, one is concerned to *remove* the object. However, this distinction doesn't seem to be genuine. When one encounters something disgusting in an open meadow, for instance, there is no urge to remove the object, but simply to avoid it. Likewise, when one encounters a dangerous snake in one's apartment, one may well respond by removing it. In both emotions, it is neutralizing the impact of the object that forms the behavioral theme. Of course, the proponent of the Motivation Thesis might respond, "it turns out that fear and disgust have the same fittingness conditions then!". But I take it this admission involves an explanatory cost, since fear and disgust do seem to have different fittingness conditions and are widely regarded as such. Immanent societal collapse is the appropriate object of fear but not disgust.

For a further example, consider awe and gratitude. Awe is fitting when directed at the sublime, and gratitude is fitting when directed at undeserved benefits offered in good will. These are different emotions with different fittingness conditions, but it is hard to see the difference between their *action types*. In both awe and gratitude, one experiences oneself as small or shrinking or vulnerable, and one wants to *celebrate* the emotion

object, to acknowledge how great it is. It might be replied that in gratitude, one doesn't merely want to celebrate the emotion object, one wants to offer *payback*. However, that desire is not universal to gratitude, but is in fact antithetical to it, since payback is a way of lifting the burden of being an undeserving beneficiary. Indeed, it is usually inappropriate—indeed, ungrateful—to offer payback to one's benefactor for the gift she has given. So, payback behaviors are often inappropriate when gratitude is called for, and thus the fittingness conditions of payback behaviors cannot explain the fittingness conditions of gratitude. Of course, in gratitude it is usually appropriate to give *something* back, not as payment but as a *symbol* of the gratitude one feels. Often a warm 'thank you' will do. And yet the same symbolic desire is present in awe (and many other emotions as well). When awed by the sea, one searches for some way to express one's awe. One wants to sing or dance or paint or take a picture or linger all morning. These are ways of acknowledging and celebrating the sublimity of the things that awe us. It seems, then, that awe and gratitude have clearly differing fittingness conditions, while their associated action type is difficult to distinguish.

Now, one might respond to these cases by appealing to more specific behaviors typical of these emotions. However, it is implausible that the specific behaviors motivated by emotions can explain an emotion's fittingness conditions. That's because, when *specific* emotional behaviors are considered, the fittingness of the emotion comes apart from the fittingness of the behavior. First off, fitting emotions can motivate unfitting behaviors. Consider anger. Often enough, situations that call for anger, and thus call for *some* retaliatory response, are situations in which the *specific* course one's anger motivates is inappropriate. One overreacts. One punishes too much or too little or too

publically or too late. A second example is disgust. One's disgust for rodents might lead one to unwittingly kill the last member of an endangered species. The disgust was appropriate, as was *some* effort to neutralize the impact of the rodent, and yet the *specific* behavior the disgust motivated was wrong. Conversely, unfitting emotions can motivate fitting behaviors. One's unfitting fear of ghosts might motivate one to avoid littering in the cemetery. One's unfitting sense of pride and achievement might motivate one to help people younger in their careers.

Accordingly, the Motivation Thesis faces a dilemma. On the one hand, the Motivation Thesis can appeal to more general action types, but in that case, the action types associated with emotions are not fine-grained enough to track the fittingness conditions of emotions. On the other hand, the Motivation Thesis can appeal to more specific actions, but then the fittingness of the actions comes apart from the fittingness of the emotion. Either way, it looks implausible that action types can explain the fittingness conditions of emotion.

By way of rejoinder, one might consider the *cognitive* patterns that are associated with emotions. Scarantino (2014) and Frijda (1986) speak of such cognitive patterns as "mental actions", including them in the behavioral profile of emotion; and Deonna and Teroni (forthcoming, n8) concede that they may need to appeal to such mental actions in order to individuate emotions. For instance, consider the impact of emotion on what one fantasizes about. In anger, one might fantasize about standing up to one's boss. In hope, one might fantasize about just how nice it will be to live in California. However, the same dilemma from above will arise in the case of fantasies. For, such fantasizing is about some course of action or other. If the course of action is described in *specific* terms, its

fittingness will be divorceable form the fittingness of the emotion associated with it. If described in *general* terms, the course of action will not be fine-grained enough to distinguish the fittingness conditions of emotions.

Alternatively, if one wants to make use of cognitive patterns, one might appeal to the impact of emotion on *attention*. In awe, for instance, one is transfixed on the emotion object.[35] However, many emotions capture attention in this way, such as fear, disgust, grief, regret, resentment, amusement, hope, pride, and loving affection. Thus, the appeal to cognitive patterns does not seem promising as a way to make fine distinctions.

I conclude, therefore, that the best unifying explanation makes use of the Evaluative Content Thesis. Without evaluative content, it is difficult to articulate a plausible explanatory relation between the fittingness conditions of an emotion and its phenomenal character.

3.5 CONCLUSION

I've argued that each emotion necessarily has its own proprietary feel, and necessarily has its own proprietary fittingness conditions. We should not accept this as a brute coincidence unless all unifying explanations are implausible. Fortunately, there are some (prima facie) plausible explanations of the coincidence. However, all plausible explanations accept that emotional experience represents evaluative properties. If emotions merely represent their propositional contents, or bodily states, or if they

---

[35] Suggested by Fabrice Teroni (in conversation).

represent nothing at all, then we cannot explain how it is that the phenomenal character and the fittingness conditions of an emotion co-vary across possible worlds.

This has implications for several other debates some of which I have mentioned at the outset. First, this directly constrains our theorizing about emotions. Second, this opens interesting prospects in the epistemology of value. Perhaps emotional experience is capable of justifying evaluative beliefs that affirm their content. Third, this is a victory for representationalists, since it establishes an explanatory relation between content and phenomenology.

Yet a further implication is this: emotional experience has *non-conceptual* content, that is, content that is represented independently of the subject having the relevant concepts. For, the Argument from Massive Coincidence reiterates in the case of children. The emotions of children can be individuated by their fittingness conditions. And presumably the emotions of children can be individuated by their phenomenal properties as well, given that adult human emotions can be. So, the Argument from Massive Coincidence will apply to the emotions of children, at ages where it is plausible that evaluative concepts are lacking. So, a child's experience of anger represents something as a wrong or an offense, prior to the child having concepts of those evaluative properties. Accordingly, those who claim that all content is conceptual must find a way to either reject the Argument from Massive Coincidence, or its extension to the case of child emotion.

4.0 EMOTIONS AS ICONS OF VALUE

It is highly intuitive that there is a deep connection between emotion and value, between fear and harm, anger and injustice, disgust and foulness, shame and shamefulness, admiration and beauty, longing and the sublime. What exactly is this connection, and in virtue of what does it obtain? Many non-realist theories of value offer ready answers. For instance, according to *sentimentalists*, our emotions or sentiments in some way construct the evaluative facts (D'Arms 2013; D'Arms and Jacobson 2000, 2006; Lewis 1989; Mulligan 1998; Prinz 2008). If sentimentalists are right, it is easy to see what the emotion-value connection is. However, many reject sentimentalism, opting instead for realist theories of value. Here, a theory of value is realist *iff* it holds that the evaluative facts obtain independently of how we represent or respond to them, including how we emotionally respond (Audi 2005; Boyd 1988; Brink 1989; Enoch 2011; Huemer 2005; Jackson 1998; Nagel 1986; Parfit 2011; Scanlon 2014; Shafer-Landau 2003).

Importantly, it is less obvious how a realist theory can account for the emotion-value connection. One option is the so-called sensibility theory (McDowell 1985; Wiggins 1976). This view provides an explicit account of the emotion-value connection: Values can be fully explained in terms of the emotions they merit, and emotions in terms

of the values that merit them.[36] However, while this view has attractions, it faces

challenges that make it reasonable to consider alternatives. First, the view embraces an

explanatory circle that many will find unattractive, although defenders of the view regard

it as benign. Second, this view faces the so-called "wrong kind of reason" problem (Crisp

2000), and although there have been several attempts to address the problem, it remains

unclear whether it can be put to rest.[37] In light of these considerations, it is worth

considering how realists might offer an alternative account of the emotion-value

connection.

Still, one might take a lead from the sensibility theory, and propose that there is

an easy way the realist can explain the emotion-value connection that silences the need

for any further theorizing: Values make emotions fitting (appropriate, reasonable, etc.).

Fear is the fitting response to impending harm, anger to injustice, etc. This is not to say,

as the sensibility theory does, that values can be *fully explained* in terms of the emotions

they make fitting. That values make emotions fitting is just one thing about them. Indeed,

this makes a fairly obvious claim about the emotion-value connection.

However, this explanation raises the further question: What is it that makes it the

case that, say, fear is the fitting response to impending harm? Why not disgust or

admiration instead? After all, plausibly this is not a brute fact. We should expect there to

be an explanation.

---

[36] Although this view makes reference to our emotional responses, it can be regarded as a realist position. Even if everyone were happy about some evil act, the act could still be an evil, as long as whether it merits happiness does not depend on how people actually respond to it.

[37] Crisp (2000) considers the example of a demon who will punish you unless you desire to swallow a cup of mud. Since this gives you good reason to desire this, such a desire seems to be merited, but clearly this is not the kind of reason the sensibility theorist had in mind. (This problem applies equally to sentimentalists such as D'Arms and Jacobson.) Thus, for those who wish to analyze values in terms of the emotions they give us reasons for (or merit, or make fitting, etc.), it is pressing to say what the right kinds of reasons are.

Fortunately, the dominant view of emotion at present affords us such an explanation. On this view, emotions are perceptual states that represent evaluative properties, and this is what explains their fittingness conditions (Prinz 2004; Roberts 2003; Tappolet 2012; Tye 2008). The fittingness of emotion is thus understood as a kind of *accuracy*: fitting emotions are those that accurately represent evaluative properties. For instance, perhaps fear perceptually represents its objects as dangerous, and that is why fear is a fitting response to the dangerous. Call this *the perceptual theory*. The perceptual theory was preceded by the view that emotions are evaluative *judgments*, but the perceptual theory has gained favor largely because of the ways emotions seem to come apart from evaluative judgments, as when you fear an insect you know to be perfectly safe.

If emotions are evaluative perceptions, then this furnishes an account of the emotion-value connection, and it even comes with an epistemological bonus: perhaps emotional experience provides evaluative justification or even knowledge on the model of sense perception. Recently, many realists have embraced the perceptual theory, or something like it, partly in light of this epistemological promise (Cowan 2015; Cuneo 2006; Dancy 2014; Döring 2003; Johnston 2001; Milona 2016; Oddie 2005; Pelser 2014; Roberts 2013; Tolhurst 1991; Wedgwood 2007). As Mark Johnston remarks:

> [W]e desire other things and other people, we are struck by their appeal, we are taken with them. This is part of how things are manifest to us: part of their appearing or presenting is their presenting to us in determinate ways and to various degrees appealing or repulsive. On the face of it, appeal is as much a manifest quality as shape, size, color and motion (Johnston 2001, p. 188).

The major difference between this view and the sensibility theory is that here, values are thought to be monadic properties out there in the world, and the role of emotional experience is simply the epistemological one of detecting them, much like the role of visual experience with respect to primary qualities.

If emotions represent values, however, then here again we should expect an explanation as to why. Plausibly, this is not a brute fact, any more than the fittingness of emotion is. After all, a fundamental project in the philosophy of mind is to explain what it is that determines the contents of mental states. It is reasonable to expect some semantic theory or other to be able to explain the evaluative content of emotion.[38] Indeed, the view that emotions are evaluative perceptions is not complete until coupled with such a semantic theory. However, I will argue that, on current semantic theories, emotions turn out to have evaluative content in a way that is extremely arbitrary (in a sense to be explained); and thus, the deep connection between emotion and value is fumbled mid-field.

Accordingly, realists face a dilemma, provided they wish to use the perceptual theory to explain the emotion-value connection. On the one hand, if we have no mental semantics for the evaluative content of emotion, then the perceptual theory remains underdeveloped. On the other hand, current semantic theories make the relation between emotion and value highly arbitrary, and so devastate the perceptual theory's ability to explain the emotion-value connection (so I will argue).

---

[38] But for a skeptical view about the project of explaining mental content, see Stich (1996), Ch. 5.

In this chapter, I will address this problem by suggesting how we can understand the semantics of emotion without treating their evaluative content as arbitrary. My central claim is that emotions are *icons of value*. Emotions represent values in part because they resemble them.[39] That is why there is a deep connection between emotion and value, and I will explain this connection in a way that is amenable even to the most robust forms of realism.

Although I will explain iconicity later on, here I offer the following loose definition: A representation is iconic *iff* in order to fully explain its content we must in some way appeal to the fact that it replicates certain features of what it represents. Importantly, note that I do not claim iconicity is either necessary or sufficient for representation, only that it sometimes partly explains why a mental state has the content it does, and that in the case of emotion, it can get us very far indeed.

In what follows, I will simply take for granted that emotions represent evaluative properties, and that they do so perceptually. I am not trying to establish the perceptual theory from scratch, only to enhance it with a semantics that can respect the emotion-value connection.

In section 4.1, I will discuss current semantic theories and show how it is they treat emotions as having an arbitrary link to the values they represent. In section 4.2, I will sketch the idea that emotions are icons of value, and show how this better respects the emotion-value connection. In section 4.3, I will bolster my view by specifying features of emotion that bear iconic relations to value.

---

[39] Nozick (1989) briefly discusses this idea (pp. 93-98), but to-date there is no version of the view brought into maturity.

4.1 THE PROBLEM WITH CURRENT MENTAL SEMANTICS

If emotions represent values, there will be some explanation to why. The most developed answer to this is due to Prinz (2004), and accordingly I will focus on his account, and show that it fails to explain an important aspect of the way emotions represent values. Importantly, this is not a quirk of Prinz' view. Since he offers a fairly direct application of informational semantics to emotions, my criticism will apply to informational semantics more generally. The upshot is that those who wish to hold the perceptual theory of emotion must find an alternative mental semantics.

4.1.1 PRINZ' SEMANTICS FOR EMOTION

Following Dretske (1981, 1995), Prinz holds that "a mental representation is a mental state that is reliably caused by something and has been set in place by learning or evolution to detect that thing." (2004, p. 54). Now, Prinz speaks of reliable causation here. Elsewhere he speaks of reliable *co-occurrence*, and I will use this term in order to respect the view that values are causally inefficacious. Accordingly, on Prinz' view, for emotions to represent evaluative properties, it must be the case that they (a) reliably co-occur with value, and (b) have the function of detecting values. For instance, fear represents threats because it reliably co-occurs with threats, and has the function of detecting them.

This much is typical information semantics, but the distinctive thing about Prinz'

view is that he claims emotions detect values *by* monitoring the subject's bodily states:

> Each emotion is both an internal body monitor and a detector of dangers, threats, losses,
> or other matters of concern. Emotions are gut reactions; they use our bodies to tell us how
> we are faring in the world (2004, p. 70).

Thus, emotions detect values by tracking changes in the body. The bodily changes in

question, of course, are not simple processes like the mere racing of the heart. Rather, the

idea is that each emotion has a unique and therefore suitably complex configuration of

bodily symptoms. Each emotion has a somatic signature. (The plausibility of this is of

course a controversy as old as the James-Lange theory itself.)

However, I will argue that this approach to the content of emotion fails to explain

something important about the way emotional experience represents values, and in what

follows I will articulate just what that is.

4.1.2 ARBITRARINESS OF REPRESENTATIONAL FORMAT

Representations can have the same content but differ in format. I can warn you the

bandits are coming by uttering this in English, but I can also use smoke signals—same

content, different format. Some formats exhibit interesting differences. Consider the

difference between linguistic and imagistic formats. The English word 'duck' refers to

ducks, but 'frog' would have done just as well. In contrast, a drawing of a duck can be

used to refer to ducks, but a drawing of a frog would not have done just as well. Of

course, you *can* use an image of a frog to refer to ducks. But then you would be treating

the image as a word and not as an image. To capture this distinction, I will say that the relation between a word and its content is *arbitrary*,[40] while the relation between an image and its content—when the image is used *as* an image—is *non-arbitrary*.

This contrast between words and images has often been noticed. For instance, as Greenberg (2013) remarks, "The relationship between a drawing, photograph, or perceptual representation of a scene and the scene itself is one of intimate correspondence, nothing like the stipulative association between a word and its denotation" (p. 216).

One might resist this. According to an influential view from the 20[th] century, pictures are governed by conventions in much the way that languages are, and so they are just as arbitrary as, say, English sentences[41] (Bryson 1983; Goodman 1968;[42] Robinson 1979; Wollheim 1970). However, following Abell (2005), we should reject this view. If convention governed picture content, then it would have to do so by governing basic elements of pictures, such as color and shape. That's because pictures are generative. We can recombine basic color and shape to depict a vast range of things. So, we have to understand convention as governing the meanings of these basic picture elements, as well as how they are combined. However, this runs into a problem when combined with a plausible view of linguistic convention. On this view, a convention is a widely adopted solution to a recurring coordination problem (Lewis 1969). Importantly, one aspect of this view is known as the *salience condition*: Those who adopt the solution must have a

---

[40] But there are two exceptions: onomatopoeia words and pictographic languages.

[41] I will often speak of specific languages, rather than language in general, since languages are extremely diverse. It seems a bit fast to assume that what we say about language, when it is logic and Latin that are foremost in our minds, will automatically apply, say, to the oracle bone script ancestral to modern Chinese (see Keightley 1996).

[42] It is tricky to classify Goodman as holding the conventionalist view, but for a defense of this classification, see Abell (2005), fn. 3.

salient conception of what the solution is (Schelling 1960). Roughly put, you can't adopt a solution if you don't know what it is and how it applies. The problem, however, is that when we consider the basic picture elements in isolation, we have no clue what they mean. For instance, in some pictures a black dot can represent an eye, but apart from the complete picture, we have no way of interpreting the meaning of the black dot. By itself, there is an indefinite number of things the dot could stand for. Thus, conventions for the meanings of basic picture elements cannot satisfy the salience condition, so any theory of convention that uses the salience condition or something like it will be incompatible with a conventionalist treatment of pictures.[43]

It remains plausible, then, that 'duck' is arbitrary in a way that the duck image is not. But while the notion of arbitrariness is evident in examples, it is difficult to define it in precise terms, and it is worth getting a grip on the notion before we proceed.

Here, convention is relevant in a different way. You might think that an arbitrary format just is one whose meaning is fully determined by convention. Certainly, this seems to be true for 'duck'.[44] However, many hold mental representation to be language-like and therefore to involve an arbitrary format, yet it is implausible, to say the least, that mental content is determined by convention . Conventions themselves depend on intentional mental states. So, conventionally determined meaning does not provide a definition of arbitrariness, only an example, albeit a paradigmatic one.

---

[43] While admittedly the salience condition has its critics (Gilbert 1989; Skyrms 1996), it nevertheless continues to attract adherents as well (Binmore and Samuelson 2006; Sugden 1986/2004, 2011).

[44] Of course, this is not the story we find in the externalist semantics of Kripke and Putnam. But there is an obvious sense in which people knew what 'water' meant even before they knew that it referred to H_2O (cf. Lycan 2006). This narrower kind of meaning is what I am discussing here.

In our example from earlier, the main difference between 'duck' and the duck image seems to be this: we can replace 'duck' with any other markings without in any way impacting the semantic content of the word. Any other markings could have just as well played the same conventionally determined role. In contrast, if we were to massively alter the image of a duck so that it more closely resembles a frog, this would seem to impact its semantic content. We can say, then, that arbitrariness is *replaceability without semantic impact*. To be more exact:

> *Arbitrariness Definition*: A format's degree of arbitrariness = the degree to which we could alter the format without impacting its semantic value.

A few remarks are in order. First, the idea of replaceability without semantic impact is probably not fundamental. More fundamental, probably, is the idea that, in part, *format determines content*. That would explain *why* we can impact semantic value by altering format. In virtue of which properties does the format partly determine content? What else must we add for full-blown representation? This gets us moving towards a substantive account of arbitrariness, and that is not our aim at present. We simply need to understand what arbitrariness is on a basic level, and how to assess the degree to which a representation has it.

Second, in our definition, it is appropriate that we have degreed notions on both sides of the "equals" sign, since arbitrariness is degreed. A highly realistic satellite picture of a terrain is less arbitrary than a map that represents the locations of cities using dots. The dots on the map could have been anything else (empty circles, Xs, etc.) without

impacting the content of the map. If, however, we dramatically altered the spatial layout

of the dots, this would impact the semantic value of the map. So, the dots are arbitrary but

the spatial layout is not. Still, on the whole the map is less arbitrary than the English

words. Of course, even realistic pictures admit of some arbitrariness. Whether someone's

hair is depicted as brown543 or brown544 will not make a semantic difference.

Third, the notion of *semantic value* here is ambiguous between two very different

interpretations: *reference* and *accuracy*. What it takes for an image to *refer* to a duck is

not the same as what it takes for an image to *accurately represent* the duck. An

inaccurate drawing of a duck may still refer to it. In what follows, I assume that some

formats are non-arbitrary in both reference and accuracy. I assume that some formats are

such that, were we to change certain features of them, they would not be inaccurate, but

rather they would not even refer to the same object. There is a related debate here as to

whether pictures must to some degree resemble their objects in order to refer to them in

the first place. Hopkins (1998) and Abell (2009) hold this view, but Greenberg (2013)

rejects it. He thinks a picture (*qua* picture) could refer to someone even if it in no way

resembled her, as long as the artist intended it to, although it would be a highly inaccurate

picture. Now, one wonders how seriously Greenberg means this. Suppose a blindfolded

artist attempted a portrait of you, but benightedly failed to make even a single marking on

the chapter. Does the blank sheet of chapter still refer to you? Plausibly, No. We would

not say, "What an inaccurate depiction". You might reply that this is because it is not a

depiction in the first place. But I don't see why not, since a blank sheet of chapter *could*

be a depiction of something, say, a fresh sheet of snow. Of course, it also more closely

resembles a fresh sheet of snow than it does a person, but that's just the point.

A further issue is that Greenberg does not consider the example of pictographic languages. Pictographs seem to refer to things at least partly in virtue of resemblance, even though a pictograph by itself makes no claim and thus cannot be regarded as accurate or inaccurate. If resemblance plays no role in explaining reference, then Greenberg must treat pictographic languages in just the same way as non-pictographic ones. Plausibly, this would mean that the reference of pictographs is just as conventionally determined as the reference of English words. In fact, however, the reasoning from Abell (2005) will repeat here. In hieroglyphic Egyptian, the term for 'flamingo' is a little image of a flamingo, and in drawing this glyph, one can use a dot for an eye (Selden 2013, p. 11), as in Figure 1:
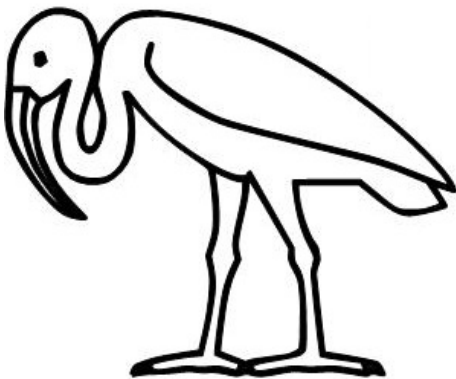

Figure 1. *Flamingo hieroglyph*. A representation of the Egyptian hieroglyph for 'flamingo'.

Here again, the dot by itself is something we have no idea how to interpret apart from the role it has in the complete image. So, while the dot is contributing to the meaning of the image, it can't be doing so in virtue of conventions, assuming we endorse the salience condition.

Now, make no mistake: convention obviously plays a significant role in hieroglyphic and other pictographic languages. They are languages, after all. For example, some ancient Egyptian texts use *cryptograms*, pictorial figures used to convey unexpected meanings that are nevertheless grounded imagistically. For instance, an image of a head can convey 'head', but can also convey 'seven'. Why? It's because a human head has seven openings (two ears, two nostrils, two eyes, and a mouth) (Selden 2013, p.12). Convention guides how to interpret which meaning, 'seven' or 'head', is conveyed, but convention does so while cooperating with the image. Convention is not playing just the same role it plays in English.

The upshot for arbitrariness is this: given that convention only partly explains pictographic reference, it is natural to think that resemblance also plays a role.[45] And to the extent that pictographs refer in virtue of resemblance, they refer in a way that is non-arbitrary. So, arbitrariness applies both to reference and accuracy.

As is clear from our examples so far, many representational systems that exhibit non-arbitrariness are those that in some way rely on resemblance.[46] The drawing of the duck resembles the duck, and that is why it is non-arbitrary with respect to ducks, while 'duck' in no interesting way resembles a duck. However, resemblance and non-arbitrariness are distinct notions. To see this, consider images that use *curvilinear projection* (or the "fisheye" perspective), as in Figure 2.

---

[45] I am aware of no philosophical work devoted to the semantics of natural pictographic languages. It is a striking omission, but one that I obviously cannot remedy here.

[46] The appeal to resemblance in explaining imagistic content is controversial (Goodman 1968), but recently there has been a resurgence of interest in the idea that resemblance plays an important role (Abell 2009; Hopkins 1998; Greenberg 2013; Kulvicki 2006; Peacocke 1987). I will discuss resemblance further when I turn to iconicity in section 4.2.

Figure 2. *Curvilinear perspective*. A picture of a railroad in curvilinear perspective.

Recently, Greenberg (2013) has argued that the accuracy of curvilinear pictures in fact

depends on the ways they do *not* resemble what they depict. However, curvilinear

pictures are nevertheless highly non-arbitrary. Even their distinctive failures of

resemblance cannot be much altered without downgrading the accuracy of the picture.

Accordingly, curvilinear pictures provide an example of non-arbitrariness without

resemblance.

Finally, whether a representation is arbitrary is not simply a matter of whether it is

easy to use. For certain purposes linguistic representations are easier to use than imagistic

ones. For instance, it is often easier to use linguistic representations when it comes to

information involving logical operations like quantification and negation. This is so, even

when the linguistic representations in question have an arbitrary relation to their contents.

For instance, even if we could develop conventions to imagistically represent *if there are*

ducks in the pond, then either it is summer time, or the zookeeper will lose his job, this

will be easier to represent using language (cf. Camp 2007, p. 168). This is so, even

though each term of the sentence is arbitrary with respect to its referent. So, arbitrariness

is not just about ease of use.

Since ease of use comes apart from arbitrariness in this way, the notion of

arbitrariness/non-arbitrariness here should not be equated with the notion of

naturalness/unnaturalness that some authors use. For instance, Giardino and Greenberg

(2015) write that

> A system is more or less NATURAL to the degree to which human nature— including relatively universal aspects of cognition, physiology, social behavior, and environmental interaction— rather than enculturation, makes that system easy to internalize and use. ... The system in which a left-hand blinking light indicates a left turn, and a right-hand blinking light indicates a right turn is especially natural in this sense: it is easy to internalize and use, presumably because it harmonizes with basic features of human cognition and body organization. The opposite system, where a left-hand light indicates a right turn is correspondingly unnatural (p. 8).

The information *if there are ducks in the pond, then either it is summer time, or the*

zookeeper will lose his job is easier to access in a linguistic representation, and it may

well be that this holds in virtue of universal aspects of human cognition. There would still

remain a sense in which the linguistic representation of this information is arbitrary

compared to the more unwieldy imagistic representation of the same. So, arbitrariness is

not naturalness in the sense above.[47]

In sum, the arbitrariness of a format is a matter of replaceability without semantic

impact. The semantic impact in question includes both reference and accuracy. If a

format is non-arbitrary, then by altering the format, we can alter its reference as well as

---

[47] Giardino and Greenberg (2015, p. 9) consider with sympathy the possibility that arbitrariness is the opposite of naturalness, where their notion of arbitrariness seems to be the same as mine, introduced with the same kinds of word-image contrasts.

its accuracy. Moreover, to say that a format is non-arbitrary is not to say that it resembles its content or that it is easy to use.

4.1.3 EMOTIONS AS NON-ARBITRARY REPRESENTATIONS OF VALUE

With the notion of arbitrariness in view, it is clear that Prinz' informational semantics treats format as arbitrary. On Prinz' view, emotions stand for values as arbitrarily as words stand for their referents. After all, consider that in principle just about anything could reliably co-occur with values. Indeed, Prinz' theory wears this result on its sleeve. Referring back to Dretske, Prinz writes:

> Consider "fuzz busters," which people place in their cars to determine when they are driving in zones monitored by police radars. A beep emitted from a fuzz buster represents the presence of a police radar. But the beep itself is utterly lacking in structure. It cannot be analyzed in to meaningful subbeeps. … [T]he beep emitted by a fuzz buster does not describe what it represents. It represents police radars because it is reliably caused by police radars, and it is set up for that purpose. Likewise, emotions can represent [values] without describing them (2004, p. 65).

The beep represents police radar just as arbitrarily as 'duck' represents ducks. Indeed, the beep could just as well have represented anything else, even the *absence* of police radar. However, while this is true of beeps, it is highly objectionable in the case of emotions. Emotions are not arbitrary beeps of value that could have just as well represented anything else.

Now, the notion of arbitrariness is most clearly on display in contrast cases between words and images. Thus, to assess whether emotions are arbitrary with respect to values, we can try to construct similar contrast cases for emotions. One such case will contrast emotional with non-emotional representations of the same evaluative property.

To illustrate, consider that there is a possible world in which olfactory sensations, rather than emotions, reliably co-occur with values, and have the function of detecting them. Imagine that in this possible world, the smell of roses—the way that roses smell in our world—co-occurs not with roses but with dangers. Such a world could also be one in which the smell of roses has the function of detecting dangers, through some bizarre twist of evolutionary fate (maybe all the dangerous things in that world happen to smell as roses smell in our world).[48]

Now, I am happy to grant that in some possible world olfactory sensations represent danger. I am not making a modal point. The point is rather that there seems to be an important difference between the way a random olfactory sensation might represent dangers in some possible world, and the way that fearful experiences represent dangers in the actual world. It seems to me intuitive that this difference is one of arbitrariness. As a representation of danger, the smell of roses is highly arbitrary in a way that the feeling of fear is not.

If this is not intuitive to you, I offer the following argument. Consider our earlier definition of arbitrariness: replaceability without semantic impact. How does this bear on the case? Notice that in the imagined case the smell of roses could have been altered in any way, or even replaced with any other smell such as the smell of magnolia blossoms, and this would not have impacted what the olfactory sensation represented. In contrast to this, the feeling of fear is not replaceable without semantic impact. For instance, one salient aspect of the feeling of fear is its aversiveness. Imagine replacing the aversiveness of fear with, say, a mild pleasant tingling in the hands. If fear involved such a pleasant

---

[48] Kenny (1963) considers a similar thought experiment.

sensation,[49] rather than aversiveness, would this impact the evaluative content of fear? Of course, I admit that any sensation, even a pleasant one, *could* represent danger, albeit in an arbitrary way, provided certain further conditions are met (whatever they are). But in our imagined case of "pleasant fear", let us suppose these further conditions are not met.

Here is a reason for thinking that to change the aversiveness of fear into pleasantness would impact the evaluative content of fear: It is generally unfitting to be pleased by dangers *qua* dangers. Of course, many people are pleased by dangers *qua* the opportunity for thrill seeking, or displaying physical prowess, or growing in courage, etc. And feeling pleasure in response to danger *qua* these things might be fitting. But it is unfitting to be pleased by danger—by potential harm—simply in itself. So, if fear were to feel pleasant, then this would affect the fittingness conditions of fear. It would no longer be fitting to fear dangers, since it is not fitting to feel pleasure in response to dangers.

One might object here that pleasures cannot be assessed for fittingness in this way. However, we evaluate pleasures all the time for fittingness as well as other normative statuses. Most of us think there are fitting and unfitting objects of sexual pleasure, for instance. There are some things one simply should not enjoy looking at. This is true of other kinds of pleasure as well. Having to kill Old Yeller is one thing; enjoying it is quite another.[50]

---

[49] Those who think emotions are identical with certain feelings might regard this as impossible. Very well. Imagine then that we are talking about a state that is like fear in all respects except that it is accompanied by this pleasant sensation in place of the aversiveness. We can even imagine such a phenomenally pleasant state as motivating fight-or-flight behaviors, since feelings are not necessary for motivation.

[50] Note that I do not assume there is a single unitary mental state that answers to 'pleasure'. Note also that it is not accurate to say these remarks beg the question against the hedonist. The problem of inappropriate pleasures is simply one of the classic problems facing hedonism (Brentano 1889/1969, p. 90; Moore 1903, sec. 56), and moreover, I do not assume the hedonist has no way of accommodating the badness of such pleasures (see Feldman 2004, Ch. 5).

The point remains, then, that to change the aversiveness of fear into pleasantness would change the fittingness conditions of fear, since it is unfitting to be pleased by danger *qua* danger. But recall that for our purposes, we are taking the perceptual theory of emotion for granted. On this theory, the fittingness conditions of an emotion are explained as a kind of accuracy condition that obtains in virtue of the emotion's evaluative content. Thus, on this view, a change in fittingness conditions would plausibly mean a change in evaluative content. Accordingly, by replacing the phenomenal aversiveness of fear with phenomenal pleasantness, we have impacted the evaluative content of fear. The feeling of fear is not replaceable without semantic impact.

Now, I can allow that from a *pragmatic* point of view, it is useful for emotions to represent the values they do. It makes good pragmatic sense for evolution to have chosen, say, fear as the detector for dangers, since fear motivates adaptive responses to dangers. As Prinz puts it, "our hearts race to increase blood flow, which prepares us for fleeing, fighting, or engaging in other kinds of behavior" (2004, p. 69). Certainly, it is interesting that emotions represent values while at the same time motivating adaptive or appropriate responses to those same values.[51] However, this does not explain our reaction to the bizarro world. Imagine a bizarro world in which cooing feelings come to represent dangers. In the bizarro world, when someone stumbles upon a deadly creature, she feels that melting sensation that we feel when a baby gives us a big smile. Now, imagine also that in the bizarro world, the forests are filled with deadly creatures that will spare your life if and only if you find them adorable. In that case, cooing feelings are a highly adaptive response towards such dangers. But there remains a sense in which to use

---

[51] Of course, emotions do not always motivate adaptive or appropriate responses, even when the emotion itself is appropriate.

cooing feelings as representations of dangers is to use them in a highly arbitrary way, while to use them as representations of the adorable is to use them in a highly non-arbitrary way.

4.1.4 SUMMING UP

In sum, we have seen that emotions have a non-arbitrary link to the values they represent. We cannot treat emotions on the model of linguistic representation. Emotions are in this way unlike the bell that indicates dinner but could have just as well indicated anything else. The bell has a totally arbitrary relation to dinner. In contrast, for emotions to represent values other than the ones they do would be to change their representational nature. Informational semantics, however, fails to respect this fact.

Informational semantics is one of the two major programs for explaining content, and the other, computational semantics, fares no better. On this view, meaning is wholly or largely derived from computational operations. This view places even fewer constraints on representational format than informational semantics does. *What* is being operated on is largely irrelevant from a semantic point of view (Block 1986; Brandom 1994; Harman 1973, 1987; Wittgenstein 1953).[52] So, abandoning informational semantics and adopting computational semantics will only worsen things. Accordingly, we have reason to seek a new semantics for the evaluative content of emotion. At the very least, we must add something to existing semantic theories in order to respect the non-arbitrary link between emotions and the values they represent. In the next section, I will offer a positive proposal as to how we might go about this.

---

[52] Note, however, that Block develops a version of the view that allows for constraints on representational format (Stich and Warfield 1994, p. 103-104).

4.2 ICONICITY TO THE RESCUE

In this section, I will introduce the notion of iconicity, and show how it gives us the resources for treating emotions as non-arbitrary representations of value.

4.2.1 ICONICITY DEFINED

To introduce iconicity, let us begin with paradigmatic instances of icons: public images, especially pictures, maps, and diagrams. These are distinct types of images, admitting of interesting differences, but what they have in common is that they represent what they do, and in the way they do, partly in virtue of *replicating* certain relevant features of the things they are images of.

Consider pictures. Plausibly, a realistic portrait of Marilyn Monroe represents her partly in virtue of replicating, in *certain* respects, what she would look like if we were to see her in person. Even Warhol's depiction of Marilyn Monroe, which massively distorts the color of her face, is a representation of Marilyn Monroe partly in virtue of replicating the shape of her face.[53]

---

[53] Here again, I take it that to be a picture *of* something is to have already achieved a certain degree of success in resembling it (Hopkins 1998). Greenberg (2013) rejects this, but I refer you to the dialectic in section 4.1.2.

Figure 3. Marilyn Dipytch. A square from Warhol's *Marilyn Diptych* (1962).

(Of course, the replication here need only be approximate.) Now, an artist might create a painting of Marilyn Monroe that involves distortions so dramatic that there is no interesting sense in which the painting replicates her features. Such a painting might still represent Marilyn Monroe, but the explanation of how it represents her would be a different one. Someone offering a theory of image content would owe us a story here, but that is not my aim. My aim is only to illustrate the notion of iconicity with respect to certain paradigmatic cases.

With the example of pictures in mind, we may define iconicity in the following way:

Iconicity Definition: a representation is iconic *iff* in order to fully explain its content, we must in some way appeal to its replicating certain relevant features of what it represents.

Notice several things about this definition. First, it is extremely schematic. I am not trying to give a substantive analysis of iconicity, only to define the notion with enough clarity to wield it. Second, this definition allows us to say, as we should, that English sentences are not iconic, since we can fully explain the contents of English sentences without appeal to any notion of replication. Third, this definition does not say that iconicity is sufficient for representation. Plausibly other conditions must be added, since many things replicate features of many other things without being representations of them or anything else for that matter. Fourth, this definition does not say that iconicity is necessary even for imagistic representation. Rothko's squares might represent something besides giant squares, but perhaps not in virtue of replicating its relevant features, and so not iconically. Sixth, this definition does not tell us exactly how replication must figure in the explanation of the icon's content. There are diverse views about this matter (Abell 2009; Hopkins 1998; Hyman 2006; Kulvicki 2006; Peacocke 1987), and there is no need to settle on one at present.

The notion of replication here can plausibly be analyzed in terms of *property sharing* (cf. Cowling 2017). A picture replicates features of the scene it represents only to the extent that the picture and the scene share certain properties, which in this case will be properties pertaining to visual appearance such as shape and color. I take it this notion of iconicity is also what authors have in mind when they speak of *resemblance*. To the extent that a picture replicates the relevant features of what it represents, the picture *resembles* what it represents.

Some iconic representations share only a small number of properties with what they represent. For example, a mercury thermometer represents iconically. The mercury rises as the temperature increases.[54] But since only this single property is shared between the thermometer and the temperature, the thermometer's *degree* of iconicity is very small.[55] This contrasts with pictures that share very many properties with what they represent in virtue of which they are the representations they are. For this reason Goodman (1968) refers to pictures as *relatively replete* compared with simple icons like mercury thermometers.

In explaining imagistic content, the use of iconicity is complex and controversial. Some authors reject the view that iconicity is explanatorily relevant (Goodman 1968; Lopes 1996, 2005). Others hold the very notion of iconicity to be suspect (Goodman 1967; Quine 1969). Obviously, it is beyond the scope of this chapter to address these issues. Accordingly, it is best to understand my claim in this chapter as provisional: assuming that the notion of iconicity is in good standing, and assuming that iconicity is explanatorily relevant to the contents of images such as realistic pictures, iconicity might also be explanatorily relevant to the evaluative content of emotional experience.

---

[54] Well, this is a bit too simple. Thermometers do not directly replicate the temperature in the way that a photograph replicates a visual scene. It is rather that there is a homomorphic mapping from the thermometer to the temperature (see Mathews 2007). Even so, the shared relation to the relevant dimension can be understood as a shared property.

[55] Of course, thermometers share many other properties with temperature, such as being important to humans. So, we need a story about how to decide which properties are relevant to iconicity. For one option, see Kulvicki (2006), Ch. 4.

4.2.2 EMOTIONS AS PICTOGRAPHIC ICONS

As we've seen, iconic representation is plausibly exhibited in the case of realistic

pictures. Yet a further case is maps. Unlike pictures, maps only need to replicate the

spatial structure of what they represent. For instance, the markers that stand for cities can

be arbitrary, as long as they are structured to replicate the relative positions of the cities

they stand for. Of course, there are many kinds of maps, from seating charts to

topographic maps, but what each of these have in common is that they use spatial

structure to represent spatial structure. In this sense, maps aim to be iconic only in their

syntax, that is, only in the way they combine their semantic elements to convey meaning

(cf. Camp 2007). In contrast, pictures aim to be iconic in both their syntax and semantics.

Pictures use iconic spatial structure to combine their semantic elements, but those

semantic elements themselves are iconic as well.

In this sense, we approach the inversion of cartographic representation by

considering pictographic languages such as hieroglyphic Egyptian, hieroglyphic Mayan,

and some West African languages such as Nsibidi (but note that these languages contain

many semantic elements that are non-pictographic as well). Modern Chinese script is an

interesting case, because it exploits iconicity in certain respects but is not

straightforwardly pictographic.[56] For instance, the Chinese character for mud is the visual

combination of the characters for water and soil. Even though the character does not have

any strict resemblance to mud, the character exploits iconicity by being formed out of

terms for the very elements that mud is formed out of. (Note, however, that only a

---

[56] The fact that Chinese is sometimes mistakenly classified as pictographic in an unrestricted way has
drawn much scorn and taken on an air of political significance (e.g., Boltz 1986, p. 406; DeFrancis 1984, p.
xi). In light of this issue, note that I describe Chinese as iconic only in specific and restricted ways.

minority of Chinese characters work in this way). Modern Chinese is also interesting because its ancestral forms tracing back to the oracle bone script are increasingly pictographic, and so we have examples of scripts that are interestingly transitional (Yong et al. 2008, pp. 18-21).

The syntactic principles of these languages are far less iconic than their semantic elements. We find pure cases of this in contrived systems of public signs, such the OSHA Hazard Communication Standard, in which simple pictures of different kinds of dangers are used to represent those very dangers (an explosion for explosive hazards, a fire for flammable hazards, etc.). Thus, while cartographic systems rely on iconic syntax and non-iconic semantics, pictographic systems rely on iconic semantics and non-iconic syntax.

If emotional experience is iconic, it is usefully compared to pictographic systems of representation. On my view, each emotion stands for a specific evaluative property, rather than the entire event, object, or state of affairs that instantiates that evaluative property. The feeling of fear stands for danger, the feeling of disgust stands for the foul, anger for injustice or offenses, awe for the sublime, amusement for the comical, etc. So, when you see a growling dog, your visual experience elicits your fear, which then attributes DANGER to the object of your visual experience, namely, the growling dog. The emotion merely supplies the evaluative property, and together with the visual experience, the growling dog is represented as a danger. Of course, it needn't be perception that elicits the emotion; it could also be belief, imagination, or memory. If, say, you come to believe by testimony that an imperceptible poisonous gas is filling the room, this will elicit fear, even though you do not perceive the gas.

This, of course, allows that emotions are about objects in the world. On the face of it, such objects include both growling dogs and propositions about growling dogs. But this object directedness is not the same as property attribution. Fear and pride can both be directed at the growling dog, but they attribute differing evaluative properties. Plausibly, the story of how an emotion acquires its object is different than the story of how it comes to attribute the evaluative properties it does. Only with respect to the latter do I claim iconicity is relevant.

The point that emotional experience is iconic on the model of pictographic systems is important to see for two reasons. First, it fits comfortably with the fact that, although emotions are iconic, they are often elicited by non-iconic states such as beliefs. Second, it allows for a variety of views about the syntax of "mentalese" (Fodor 1975). On my view, the claim that emotions are iconic is entirely compatible with the language of thought hypothesis; it only requires that some of the semantic elements of the language of thought are pictographic in nature.[57]

4.2.3 ICONICITY IN THE MIND

So far, we have considered maps, pictures, and other public artifacts that sometimes represent iconically. It is useful to see, however, that independently of emotions, there are plausible—though controversial—instances of mental states that represent iconically.

---

[57] This contrasts with recent discussion concerning the extent to which the language of thought could be map-like. Camp (2007) argues that some thinking occurs with "mental maps", and that these are not language-like, on account of their essentially spatial syntax, while Blumson (2012) argues that maps and languages in fact use the same syntactic principles, and thus a map-like architecture is not a genuine alternative to the language of thought. In contrast, I am saying nothing about the syntax of thought, only that some of the semantic elements are icons, whatever the syntax is that combines them.

This is important to see because the claim that emotions are icons of value is much less surprising if a range of other mental states also exhibit iconicity.

The first example is *analogue magnitude representation* (AMR). AMRs offer a plausible explanation of the ways in which some animals can assess magnitudes without possessing numeric concepts (Beck 2015). For instance, ducks are capable of making decisions about food-seeking that seem to factor in the rate of food distribution and the number of food competitors (Harper 1982). Importantly, AMRs would be iconic representations since they would proportionately replicate the magnitudes they represent.

Second, some authors view the content of perception as iconic. For instance, Tyler Burge remarks that "the elements in visual perception have something like the form of a map or sketch from an egocentric perspective" (2010, p. 540). Similarly, Ned Block considers the format of perception to be "map-like or iconic" (2014, p. 560).

Third, iconicity is also apparent in hallucinatory perception, and in such cases it is plausible that visual iconicity is phenomenal. In certain hallucinations, what one experiences is image-like. For instance, Oliver Sacks describes the hallucinatory experiences of one of his patients with Charles Bonnett Syndrome who

> saw little people a few inches high, like elves or fairies, with little green caps, climbing up the sides of her wheelchair. There were children, too, "picking up pieces of chapter from the floor," or climbing hallucinatory stairs in one corner of the room (Sacks 2012, p. 21).

Taking such reports at face value, this hallucinatory experience seems to be iconic, and iconic in virtue of its phenomenal character.[58]

Now, elves and fairies do not really exist, and so it is not quite right to say that the patient's hallucinatory experience replicates the relevant features of elves and fairies. Indeed, there is a well known problem here concerning how to allow for *mis*representation within iconic systems. For instance, one might want to say that the patient's experience depicts what elves and fairies *would* look like if they did exist (Hyman 2006). But this proposal is not entirely satisfactory, since, for starters, iconic representations of non-existent objects can still misrepresent those objects, and in such cases it will be false that the icon depicts what the object would look like were it to exist (Abell 2009, p. 189). However, it is reasonable to set this complication aside. Since there are a range of puzzles raised by fictional entities and other nonexistent objects, this does not seem to be a special problem about iconicity (cf. Greenberg 2013, p. 220).

As a final example, some hold that *mental imagery* is an iconic form of representation. Here, mental imagery is sometimes defined as a form of stable (though non-hallucinatory) perceptual phenomenology that persists in the absence of stimuli (Kosslyn et al. 2006, p. 3-4). Are roses a darker shade of red than ladybugs? For proponents of mental imagery, we can consult mental images to answer such questions. On the other hand, mental imagery is sometimes understood as a form of *non*-conscious representation that is somehow picture-like (cf. Block 1983, p. 506-507). Now, the existence and nature of mental imagery is fraught with controversy, and some theorists opt to explain all mental representation in terms of more language-like formats (Pylyshyn

---

[58] I do not assume that genuine perception should be understood in the same way as hallucinations. For all I've said, the disjunctivist treatment of hallucination is correct (Hinton 1973; McDowell 2008; Snowdon 1979).

1973, 1981, 2003). This issue is highly controversial, followed by a vast literature since the 1970s. My point here is not to weigh in, but to note that if there is mental imagery, then we have yet another form of iconic representation.

The purpose of considering these examples is to see that, in addition to public artifacts, there are also mental representations that can reasonably be considered iconic. Obviously, nothing I've said compels you to take an iconic view of these mental states. The point is rather that some reasonable people do take this view, and this lends some further credibility to the idea that emotional experience represents iconically.

4.2.4 SUMMING UP: THE PROMISE OF ICONICITY

In sum, iconic representations are those whose content is partly explained by the way they replicate certain features of what they represent. A range of public artifacts exemplify this, such as pictures and maps. But emotions are best thought of on the model of pictographic systems: plausibly each emotion is iconic only with respect to specific evaluative properties, which allows that emotions can be employed in a language of thought whose syntax and other semantic elements are non-iconic. Emotions are iconic predicates, rather than a full-fledged system of iconic representation. In addition to public artifacts, moreover, many hold there are a range of mental representations that are iconic. Thus, the view that emotions are icons of value regards them as part of this broader class.

Our ultimate aim, recall, is to explain why emotions have a non-arbitrary link with value, and to do so in a way that is amenable to value realism. Iconicity has great promise here. Indeed, the hallmark feature of iconic formats is their non-arbitrary link to what they represent (Kosslyn et al. 2006, p. 12; Giardino and Greenberg 2015, p. 2-4).

Icons represent partly in virtue of the properties of the format itself, and so we can see why altering the format might impact its semantic value. Accordingly, iconicity is exactly the sort of notion we need in order to explain the non-arbitrary link between emotions and the values they represent. Thus, we have good reason to seek a view on which emotions are icons of value.

If emotional experience is iconic with respect to value, several further issues need to be addressed. My aim here has been only to introduce this view and show that it is well-motivated. Still, the most pressing question is, What features do emotions share with values such that by sharing those features, they are iconic with respect to values? A robust answer to this question is crucial for fully developing the view, and this demands a chapter of its own. In the next section, however, I will explore some options I think are plausible, in order to display the initial promise of the view.

4.3 THE ICONICITY OF EMOTIONAL EXPERIENCE

In this section, I will highlight features of emotions that *partly* explain why emotions are iconic with respect to the values they represent. Although a full account of emotional iconicity is too ambitious at present, it is important to see that there are some plausible candidates concerning the properties of emotions that can explain their iconicity. Such properties I will call *Icon Properties*, or *ICPs* for short. Plausibly, there are four conditions on identifying ICPs for emotions:

1. *Commonality*: An ICP must be a property that emotions share with the values they represent.

2. *Efficacy*: An ICP must be a property that could plausibly contribute to an explanation as to why an emotion represents the value it does.

3. *Regularity*: An ICP must be a property located in a property space such that all emotions have a property located in that space.

4. *Specificity*: ICPs, taken together, must be such that no two emotions are the same in all their ICPs but represent a different evaluative property.

Each of these conditions is intuitively plausible. Commonality and Efficacy fall out of the Iconicity Definition, provided that replication can be analyzed in terms of property sharing. Regularity is needed in order to achieve an account of emotional content that is suitably principled. And Specificity follows from the fact that (i) we are looking for properties that are sufficient for emotions to be icons of value, together with the plausible claim that (ii) each emotion represents a different evaluative property.

In what follows, I will argue that two plausible ICPs are *valence* and *intensity*.

4.3.1 VALENCE AND INTENSITY

One feature of emotional phenomenology is that it can be described in terms of valence and intensity (or arousal) (Russell 1980; Kuppens et al. 2013). Here, valence is simply pleasantness or unpleasantness. For instance, there is clearly something unpleasant about feeling depressed, anxious, angry, or bored, while there is clearly something pleasant about feeling joyful, proud, amused, or elated.

Valence is cross cut by intensity. Terror and elation differ in valence, yet both are *high* in intensity. Contentment and boredom differ in valence, yet both are *low* in intensity. Here, 'arousal' would do just as well as 'intensity', and is indeed the more common term, but 'intensity' is useful for emphasizing the felt character of arousal. For further examples of how emotions can be assessed for valence and intensity, see Figure 4.
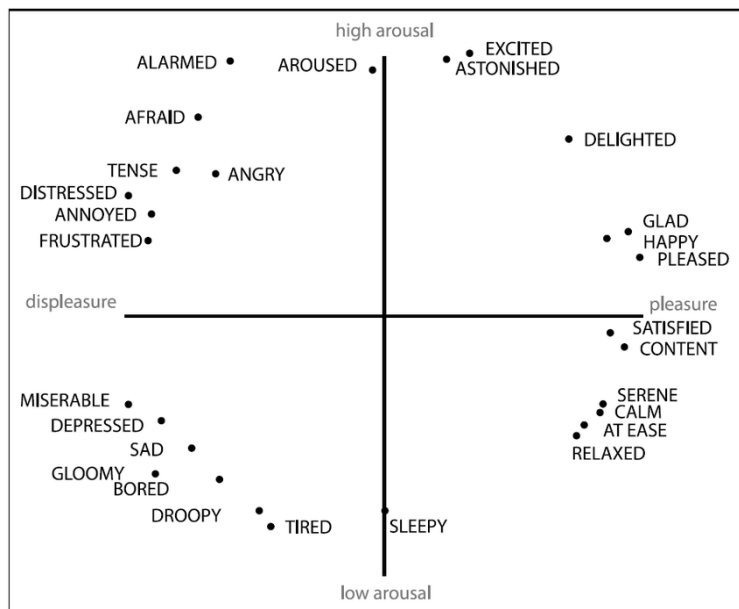

Figure 4. Affective circumplex. A chart representing various affective and emotional states along the dimensions of valence/pleasure and intensity/arousal. Adapted from Russell (1980).

The valence and intensity dimensions of emotion have been the rallying point for so-called constructionist theories of emotion (Russell and Barrett 1999), which claim that emotions are psychological constructions rather than natural kinds. However, I take no stand on this issue. It is only the notions of valence and intensity I want to employ; and obviously, emotions can be natural kinds and still be describable in terms of valence and intensity.

One issue here is how valence and intensity are related. A natural idea is that they are independent dimensions. However, the balance of evidence at present seems to favor that they exhibit a V-shaped relationship (Kuppens et al. 2013), as seen in option (d) from figure 4. The idea in (d) seems to be that intensity just is the degree of pleasantness or unpleasantness an emotional experience exhibits. This predicts that there will not be emotions that are, say, high in intensity but only very mild in pleasantness or unpleasantness, a prediction that bears out upon reflection, with a potential exception being surprise.
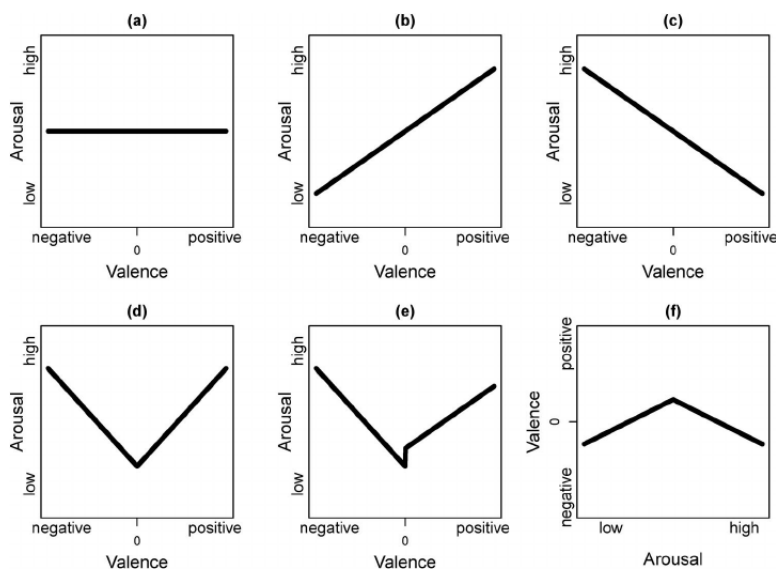

Figure 5. *Valence-arousal relations*. Six graphs representing possible relations between valence and arousal/intensity. Adapted from Kuppens et al. (2013), which argues that the balance of evidence favors (d).

4.3.2 VALENCE AND INTENSITY AS ICPS

How do valence and intensity fare with respect to the four conditions above? Valence appears to satisfy Commonality, since values also can be assessed as positive or negative. The prediction here is that negatively valenced emotions will represent negative evaluative properties, and positively valenced ones will represent positive evaluative properties. This prediction seems exactly correct, if you consider each emotion (for instance, see Figure 4).

Intensity, too, satisfies Commonality. Consider that values admit of degrees. The Gowanus Canal is foul but not as foul as an open river of sewage. Since values admit of degrees, intensity can be understood as tracking this. So, although the broad category of fear represents things as dangers, it's plausible that, say, mild anxiety represents lesser dangers, while sheer terror represents grave dangers. This is especially plausible given that the fittingness conditions of an emotion take into account the emotion's felt intensity. It is unfitting to be enraged by something that is only mildly offensive, though some lesser degree of anger might have been appropriate. This is just what we should expect, if felt intensity partly explains the evaluative content of an emotion, and the evaluative content of an emotion explains its fittingness conditions.

As for Efficacy, it is plausible that valence and intensity can contribute to an explanation of emotional iconicity, since, as noted above, fittingness conditions take into account valence and intensity in a way that is rather finely grained.

As for Regularity, it seems that every emotion can be assessed along the valence dimension. Perhaps some emotions, such as surprise, are neither positive nor negative, but neutral. Still, if surprise is an emotion, it can be located on the valence dimension by

placing it in the middle. Accordingly, valence satisfies Regularity. The same holds for intensity, since every emotion seems to land somewhere on the intensity dimension.

Finally, as for Specificity, valence and intensity by themselves fail this test. For instance, consider that terror, rage, and despair are each negatively valenced, and each can be comparably intense. But plausibly they have differing evaluative contents, since their fittingness conditions differ. For instance, rage is fitting when directed at grave offenses, while terror fitting when directed at grave dangers.

This shows that valence and intensity by themselves are not sufficient as ICPs. However, it is nevertheless plausible that valence and intensity can be included alongside other ICPs, in light of the considerations above. After all, valence and intensity satisfy the first three conditions, and though valence and intensity are not exactly fine grained enough, they do closely track the fittingness conditions of emotions. Fortunately, moreover, it is also true that terror, rage, and despair *feel* differently. They share their phenomenal character with respect to valence and intensity, but there is clearly a sense in which terror feels differently than rage does. So, just as valence and intensity are not fine grained enough to track the fittingness conditions of emotion, they are also not fine grained enough to fully characterize emotional phenomenology. This means there are aspects of emotional phenomenology that we might yet appeal to as ICPs.

4.3.3 SUMMING UP

In sum, emotional experience can be characterized in terms of valence and intensity, and these phenomenal qualities bear iconic relations to values, and thus begin to furnish us with an account of the iconicity of emotion, although the account needs further developing, especially with respect to locating a sufficient set of ICPs.

4.4 CONCLUSION

We began by noting that there is a deep connection between emotion and value. Sentimentalists have a ready explanation of this, but such an explanation is harder to achieve for realists. What I have offered is a way for realists to explain the emotion-value connection. Along the way, I have also offered the beginnings of a novel semantic theory for the evaluative content of emotion, a semantic theory with significant advantages over current alternatives. My view is that emotions are icons of value. Emotions represent values by replicating crucial features of those values. This view respects the non-arbitrary link between emotion and value, and is compatible with robust forms of value realism.

The crucial next step for this theory is to identify additional features of emotional experience that explain its iconicity. As a down payment, I've highlighted valence and intensity, but as we've seen they are not sufficient, although they will plausibly form part of any final account.

5.0 SUMMARY AND CONCLUSIONS

The guiding concern of this dissertation has been to characterize the place of emotional experience in the domain of value, and to do so in a way that (i) is compatible with robust forms of value realism, and (ii) gives emotional experience a significant and favorable role. While the three chapters can be read independently, they progressively defend a unified picture, summarized in the following theses:

1. Emotions provide epistemic benefits with respect to value. (Chapter 2)

2. No account at present has succeeded in saying what those benefits are and how emotions provide them. (Chapter 2)

3. In order to characterize the epistemic benefits of emotion, it will be useful to first establish that emotions are representations of value. (I did not claim this explicitly, but this links Chapters 2 and 3.)

4. While many authors hold the fittingness conditions of emotion directly establish that emotions are representations of value, this is mistaken. (Chapter 3)

5. But the fittingness conditions of emotion are relevant in the following way. Each emotion has a unique phenomenal character, and is governed by a unique set of fittingness conditions. This co-variance between phenomenology and fittingness is best explained by saying that emotions are representations of value. (Chapter 3)

6. Emotions represent values in a distinctive way: Emotional experience is highly non-arbitrary with respect to the values it represents. But current theories of content determination fail to respect this fact, treating emotions as arbitrary in the way that conventional linguistic representations are. (Chapter 4)

7. We can do better by modeling the representational nature of emotion after non-linguistic forms of representation, such as pictorial and cartographic representations. Emotions represent values in part because they resemble them. (Chapter 4)

There are, of course, many further issues to address. The two most pressing are the following. First, in light of thesis 7, my account must specify in virtue of which features emotional experience bears a resemblance with the values it represents. Second, in light of theses 1 and 2, my account must offer a positive view of the distinctive epistemic benefit that emotional experience provides. In what follows, I want to suggest how we might approach the latter. To address this, we must answer the following two questions:

The Grounding Question: In virtue of what do emotions provide the epistemic benefit they do?

The Value Question: What is the nature of the epistemic benefit that emotions provide?

The idea that emotions are icons of value offers promising materials for addressing these questions. Let's start with the Grounding Question. Recall, for instance, those theorists, such as Roberts and Dancy, who wanted to say that the epistemic benefits of emotion result from its being a form of appearance or seeming (Section 2.3.2). The problem, as we saw, is that there can be purely intellectual states that are also evaluative seemings, but these states do not confer the same epistemic benefits. Thus, we must appeal to a property that emotions have but that merely intellectual states do not. Iconicity is plausibly such a property. For instance, we saw that the rich phenomenological features of emotions, such as valence and intensity, are relevant to their iconicity (Section 4.3). But cold intuitions lack any such phenomenology. (They may co-occur with such phenomenology, but the connection is not typically regarded as essential.) So, it is difficult to see how intuitions could exhibit any interesting degree of iconicity with respect to values, while this is much more plausible in the case of emotional experience, in light of its rich phenomenal character.

Moreover, it has been widely noted that iconic representations have different epistemic properties than language-like representations. The adage "a picture is worth a thousand words" in fact conveys something important about iconic formats. Camp (2007) notes that pictures convey information with "rich, multi-dimensional specificity" (p. 156), and that maps and diagrams can provide more expedient ways of meeting certain cognitive demands (p. 161). Dretske (1981) notes that photographs are more informationally dense than verbal descriptions (p. 137). And many have noted that the non-arbitrary relation between icon and content makes iconic systems easier to learn and use (Giardino and Greenberg 2015, p. 8; Bordwell 2008, p. 61-63). On the face of it,

then, icons have a unique epistemic profile, and so the idea that emotions are icons of value is plausibly relevant to why they provide distinctive epistemic benefits.

I do not claim, of course, that any iconic representation automatically confers epistemic benefits. For instance, I see no epistemic benefit arising from using a mercury thermometer rather than a digital one. However, on my account, the iconic relation between emotion and value is much tighter, much more profound and intimate than the relation between a mercury thermometer and the temperature. At least, the view that emotions are icons of value is most compelling if we can develop it in this way.

Then how exactly should we answer the Value Question? One promising idea is that emotional experience *acquaints* us with values. As Roberts (2013) put it, emotional experience provides "more intimate cognitive contact" with values (p. 52), an epistemic improvement which can be regarded as a form of "personal acquaintance" (pp. 39-40).

Now, acquaintance is usually understood to be an *unmediated direct cognitive access to objects*.[59] However, the idea that emotions provide this direct form of acquaintance faces at least three challenges. First, it is tempting to think acquaintance requires some kind of causal efficacy on the part of the objects of acquaintance. It seems that the objects of acquaintance, even mental states like pain, would need to be able to cause or at least explain our awareness of them. However, many realists do not want to say that values are causally or explanatorily efficacious, and the present work seeks to accommodate all forms of realism. Second, emotions typically depend on other mental

---

[59] Bertrand Russell's definition of acquaintance: "We shall say that we have *acquaintance* with anything of which we are directly aware, without the intermediary of any process of inference or any knowledge of truths" (1912, p. 78; see also 1910/11, 1914). Contemporary uses of the notion of acquaintance can be found in BonJour (2001), Conee (1994), Fales (1996), Fumerton (2005), Gertler (2012), and Johnston (2004).

states in a way that seems to mediate the cognitive access they provide. When you see the wild horses, your longing in part depends on your visual processing of the wild horses. So, there seems to be something standing between you and the beauty. (Note that I say visual *processing*, not visual *experience*, in order to avoid the suggestion that experience is a thing standing between you and the horse.)

This point is related to the third challenge: Emotions can be elicited by mental states that do not require the presence of their objects, such as imagination, belief, and memory. When in your old age you remember the wild horse, and your longing returns, it is harder to see how you could be experiencing direct unmediated cognitive access to their beauty. Even worse, consider emotions in response to fictions and imaginative states. You may find it delightful to imagine a lazy afternoon in the Shire, but can you really be directly acquainted with the value of a thing that does not exist?

So, while I agree with Roberts that emotional experience acquaints us with values in *some* way, plausibly it doesn't do this in the direct way that acquaintance is usually understood. Accordingly, perhaps instead emotional experiences provide a kind of acquaintance that is *indirect* and *mediated*. This may strike one as an absurd proposal. Usually, saying that acquaintance is direct and unmediated is the very way the concept is introduced! However, consider that familiar examples make vivid a form of acquaintance that is not direct in the traditional Russellian sense. Consider the difference between reading a biography of Churchill, and meeting the man himself. When you meet him in person, perhaps you don't have direct unmediated cognitive access to him in the Russellian sense (of course, a direct realist will hold that you do). However, there is clearly also a sense in which meeting him in person gives you *far more* direct access to

him than reading about him does, no matter how thorough the biography. It is correct to say that you are acquainted with Churchill, even if the acquaintance is not utterly direct. Moreover, it is plausible that this acquaintance is a better epistemic status (all else being equal).

Icons, too, seem to provide a kind of indirect acquaintance. Seeing a picture of Churchill acquaints you with his physical appearance in a way that reading about him cannot. That's because the picture itself replicates the appearance of Churchill. Since icons replicate features of their contents, icons can indirectly acquaint us with their contents by acquainting us with the features they replicate. Imagine, for instance, a plaster model of a city you've never visited, say, ancient Alexandria. Imagine the model is highly detailed and realistic, and is displayed on a large table. You walk around the model, peering into its nooks and crannies, studying its vivid colors and spatial details. (Suppose the sculptor got all of this right.) Your friend, however, decides to stay home instead of seeing the model. But she spends the day reading a book that happens to report all and only the propositional information you acquired during your visit. She finishes the book, and forms the same beliefs you do about Alexandria's physical characteristics. It seems plausible that, by being acquainted with Alexandria *the model*, you are acquainted with Alexandria *the city* in a way that your friend is not.

If you deny this, then imagine a further case. You travel back in time and explore the real Alexandria. Clearly now you are better acquainted with the city than your friend is. One reason is that you are now acquainted with Alexandria's visual characteristics, its colors, shapes, textures, and spatial layout. You are acquainted with the city by being acquainted with its properties. Well, we are imagining that the model of Alexandria

replicates all of these visual characteristics. They are not the same token characteristics (obviously), but they are copies, and so the model provides acquaintance by proxy.

More needs to be said here, of course, but the foregoing lends initial credibility to the idea that icons provide indirect acquaintance.

Plausibly, acquaintance is an epistemic good, and one that is not reducible to propositional knowledge. After all, we can destroy your knowledge but your acquaintance survives. Imagine that when you time traveled a moment ago, you did so by choosing one mystery door out of a zillion. All the other doors would have led you to a situation in which you non-veridically hallucinate a city that is actually nothing like Alexandria, though you would be none the wiser. Luckily, you just happened to choose the only door that could lead you right. You enter the real city, and form all kinds of true beliefs about it. These beliefs *would* have been knowledge, but their truth is far too accidental. Your *acquaintance* with the city remains, however, and you are epistemically better off for having been so acquainted.[60]

If this brief sketch is correct, then we have the materials to answer the Grounding and the Value Question. In virtue of its iconicity, emotional experience indirectly acquaints us with value by directly acquainting us with the features of value that it replicates. Since emotional experience has the power to do this in virtue of its iconicity, plausibly, cold intuition cannot offer such acquaintance. Temple Grandin intellectually grasped the grandeur of the night sky, but it is we who, in being moved by it, are acquainted with its grandeur.

---

[60] A similar kind of epistemic luck is sometimes used to argue that *understanding* is distinct from knowledge (Pritchard 2009).

BIBLIOGRAPHY

Abell, Catharine, (2009), "Canny Resemblance", *Philosophical Review* 118 (2): 183-223.

Abell, Catharine, (2005), "Against Depictive Conventionalism", *American Philosophical Quarterly* 42 (3): 185-197.

Allen, Keith, (2017), *A Naïve Realist Theory of Color*, Oxford University Press.

Allen, Keith, (2010), "In Defence of Natural Daylight", *Pacific Philosophical Quarterly* 91: 1-18.

Alston, William P., (1993), *The Reliability of Sense Perception*, Cornell University Press.

Alston, William P., (1989), "A 'Doxastic Practice' Approach to Epistemology", *Knowledge and Skepticism*, ed. Marjorie Clay and Keith Lehrer, Westview Press.

Audi, Robert, (2013), *Moral Perception*, Princeton University Press.

Audi, Robert, (2005), *The Good in the Right*, Princeton University Press.

Bagnoli, Carla, (2011), "Moral Perception and Knowledge by Principles", *The New Intuitionism*, ed. Jill Hernandez, London: Continuum.

Bayne, Tim, (2011), "The Sense of Agency", ed. Fiona Macpherson, *The Senses*, Oxford University Press.

Bealer, George, (1998), "Intuition and the Autonomy of Philosophy", *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, eds. Michael R. Depaul and William Ramsey, Lanham, Maryland: Rowman and Littlefield Publishers, Inc.

Beck, Jacob, (2015), "Analogue Magnitude Representations: A Philosophical Introduction", *The British Journal for the Philosophy of Science* 66 (4): 829-855.

Berker, Selim, (2009), "The Normative Insignificance of Neuroscience", *Philosophy and Public Affairs* 37 (4): 293-329.

Binmore, Ken, and Samuelson, Larry, (2006), "The Evolution of Focal Points," *Games and Economic Behavior* 55: 21-42.

Blackburn, Simon, (2001), "Reason, Virtue, and Knowledge", *Virtue Epistemology*, eds. Abrol Fairweather & Linda Zagzebski, Oxford University Press.

Blackburn, Simon, (1998), *Ruling Passions*, Oxford: Clarendon Press.

Blackburn, Simon, (1993), *Essays in Quasi-Realism*, Oxford University Press.

Blackburn, Simon, (1981), "Reply: Rule-Following and Moral Realism", *Wittgenstein: To Follow a Rule*, eds. S. Holtzman and Christopher M. Leich, Routledge.

Block, Ned, (2014), "Seeing-As in the Light of Vision Science", *Philosophy and Phenomenological Research* 89 (1): 560-572.

Block, Ned, (1995), "On a Confusion about a Function of Consciousness", *Brain and Behavioral Sciences* 18 (2): 227-247.

Block, Ned, (1986), "Advertisement for a Semantics for Psychology", *Midwest Studies in Philosophy* 10 (1): 615-78.

Block, Ned, (1983), "Mental Pictures and Cognitive Science", *Philosophical Review* 92 (4): 499-542.

Blumson, Ben, (2012), "Mental Maps", *Philosophy and Phenomenological Research* 85 (2): 413-434.

Boltz, William, (1986), "Review of *The Chinese Language: Fact and Fantasy* by John DeFrancis", *Journal of the American Oriental Society* 106 (2): 405-7.

BonJour, Laurence, (2001), "Toward a Defense of Empirical Foundationalism", *Resurrecting Old-Fashioned Foundationalism*, ed. Michael DePaul, Lanham, MD: Rowman & Littlefield.

Bordwell, David, (2008), *Poetics of Cinema*, New York: Routledge.

Boyd, Richard, (1988), "How to Be a Moral Realist", *Essays on Moral Realism*, ed. Geoffrey Sayre-McCord, Cornell University Press.

Brady, Michael S., (2014), *Emotional Insight*, Oxford University Press.

Brady, Michael S., (2011), "Emotions, Perceptions, and Reasons", *Morality and the Emotions*, ed. Carla Bagnoli, Oxford University Press.

Brady, Michael S., (2009), "The Irrationality of Recalcitrant Emotions", *Philosophical Studies* 145 (3): 413-430.

Brady, Michael S., (2007), "Recalcitrant Emotions and Visual Illusions", *American Philosophical Quarterly* 44 (3): 273-284.

Brandom, Robert, (1994), *Making It Explicit*, Harvard University Press.

Brentano, Franz, (1874/1973), *Psychology From An Empirical Standpoint*, English language edition, ed. L. McAlister, London: Routledge and Kegan Paul.

Brink, David, (1989), *Moral Realism and the Foundations of Ethics*, Cambridge University Press.

Brogaard, Berit, and Chudnoff, Elijah, (2016), "Against Emotional Dogmatism", *Philosophical Issues* 26 (1): 59-77.

Bryson, Norman, (1983), *Vision and Painting*, Yale University Press.

Burge, Tyler, (2010), *Origins of Objectivity*, Oxford University Press.

Byrne, Alex, (2001), "Intentionalism Defended", *Philosophical Review* 110: 199-240.

Byrne, Alex, and Hilbert, David, (2003), "Color Realism and Color Science", *Behavioral and Brain Sciences* 26 (1): 3-21.

Camp, Elisabeth, (2007), "Thinking with Maps", *Philosophical Perspectives* 21 (1): 145-182.

Campbell, John, (2005), "Transparency vs Revelation in Color Perception", *Philosophical Topics* 33: 105-115.

Campbell, John, (1994), "A Simple View of Color", *Reality, Representation, and Projection*, eds. John J. Haldane and Crispin Wright, Oxford University Press.

Cannon, W. B., (1927), "The James-Lange Theory of Emotions: A Critical Examination and an Alternative Theory", *American Journal of Psychology 39*: 106-124.

Carruthers, Peter, (2011), *The Opacity of Mind*, Oxford University Press.

Chalmers, David, (2006), "Perception and the Fall from Eden", eds. Tamar Gendler and John Hawthorne, *Perceptual Experience*, Oxford University Press.

Chalmers, David, (2004), "The Representational Character of Experience", *The Future of Philosophy*, ed. Brian Leiter, Oxford University Press.

Christie, I., and Friedman, B., (2004), "Autonomic Specificity of Discrete Emotion and Dimensions of Affective Space a Multivariate Approach", *International Journal of Psychophysiology* 51: 143-153.

Chudnoff, Elijah, (2013), *Intuition*, Oxford University Press.

Churchland, Paul, (1985), "Reduction, Qualia, and the Direct Introspection of Brain States", *Journal of Philosophy* 82: 8-28.

Cicero, (2002), *Cicero on the Emotions: Tuscalan Disputations 3 and 4*, trans. Margaret Graver, Chicago Unversity Press.

Clore, G. L., Gasper, K., and Garvin, E., (2001), "Affect as Information", *Handbook of Affect and Social Cognition*, ed. J. P. Forgas, Mahwah, NJ: Lawrence Erlbaum Associates.

Conee, Earl, and Feldman, Richard, (2004), *Evidentialism*, Oxford University Press.

Conee, Earl, (1994), "Phenomenal Knowledge", *Australasian Journal of Philosophy* 72 (2):136-150.

Cowan, Robert, (2015), "Perceptual Intuitionism", *Philosophy and Phenomenological Research* 90 (1):164-193.

Cowling, Sam, (2017), "Resemblance", *Philosophy Compass* 12 (4): 1-11.

Crane, Tim, (1998), "Intentionality as the Mark of the Mental", *Contemporary Issues in the Philosophy of Mind*, ed. Anthony O'Hear, Cambridge University Press.

Crisp, Roger, (2005), "Value, Reasons and the Structure of Justification: How to Avoid Passing the Buck," *Analysis* 65 (1): 80-85.

Cuneo, Terence, (2006), "Signs of Value: Reid on the Evidential Role of Feelings in Moral Judgment", *The British Journal for the History of Philosophy* 11: 69-91.

Damasio, Antonio R., (1994), *Descartes Error: Emotion, Reason, and the Human Brain*, New York: G. P. Putnam's Sons.

Dancy, Jonathan, (2014), "Intuition and Emotion", *Ethics* 124 (4): 787-812.

Dancy, Jonathan, (1986), "Two Conceptions of Moral Realism", *Aristotelian Society Supplementary Volume* 60 (1): 167-205.

D'Arms, Justin, (2013), "Value and the Regulation of the Sentiments", *Philosophical Studies* 163 (1): 3-13.

D'Arms, Justin, and Jacobson, Daniel, (2006), "Anthropocentric Constraints on Human Value", *Oxford Studies in Metaethics*, Vol. 1, ed. Russ Shafer-Landau, Oxford: Clarendon Press.

D'Arms, Justin, and Daniel Jacobson, (2003), "The Significance of Recalcitrant Emotions (or, Anti-Quasijudgmentalism)", ed. Anthony Hatzimoysis, *Philosophy and the Emotions*, Cambridge University Press.

D'Arms, Justin, and Jacobson, Daniel, (2000), "Sentiment and Value", *Ethics* 110 (4): 722-748.

DeFrances, John, (1984), *The Chinese Language: Fact and Fantasy*, Honolulu: The University of Hawaii Press.

Deonna, Julien, and Fabrice Teroni, (forthcoming), "Getting Bodily Feelings into Emotional Experience in the Right Way", *Emotion Review*.

Deonna, Julien, and Fabrice Teroni, (2015), "Emotions as Attitudes", *Dialectica* 69 (3): 293-311.

Deonna, Julien, and Fabrice Teroni, (2014), "In What Sense Are Emotions Evaluations?", eds. Cain Todd and Sabine Roeser, *Emotion and Value*, Oxford University Press.

Deonna, Julien, and Fabrice Teroni, (2012), *The Emotions*, Routledge: New York, NY.

Dirac, Paul, (1930), "The Proton", *Nature* 126: 605-606.

Döring, Sabine A., (2015), "What's Wrong With Recalcitrant Emotions? From Irrationality to Challenge of Agential Identity", *Dialectica* 69 (3): 381-402.

Döring, Sabine A., (2009), "The Logic of Emotional Experience: Noninferentiality and the Problem of Conflict Without Contradiction", *Emotion Review* 1 (3): 240-247.

Döring, Sabine A., (2003), "Explaining Action by Emotion", *Philosophical Quarterly* 53 (211): 214-230.

Dreier, James, (2004), "Meta-ethics and the Problem of Creeping Minimalism", *Philosophical Perspectives* 18 (1): 23-44.

Dretske, Fred, (1995), *Naturalizing the Mind*, MIT Press.

Dretske, Fred, (1981), *Knowledge and the Flow of Information*, MIT Press.

Enoch, David, (2011), *Taking Morality Seriously*, Oxford University Press.

Fales, Evan, (1996), *A Defense of the Given*, Lanham, MD: Rowman & Littlefield.

Feldman, Fred, (2004), *Pleasure and the Good Life*, Oxford: Clarendon Press.

Fine, Kit, (2001), "The Question of Realism", *Philosophers' Imprint* 1 (1): 1-30.

Fodor, Jerry, (1984), "Observation Reconsidered", *Philosophy of Science* 51: 23-43.

Fodor, Jerry, (1983), *The Modularity of Mind*, MIT Press.

Foot, Philippa, (1978a), "Hume's Theory of Moral Judgment", *Virtue and Vices and Other Essays in Moral Philosophy*, Berkeley: University of California Press.

Foot, Philippa, (1978b), "Moral Beliefs", *Virtue and Vices and Other Essays in Moral Philosophy*, Berkeley: University of California Press.

Friedman, Michael, (1974), "Explanation and Scientific Understanding", *Journal of Philosophy* 71: 5-19.

Frijda, Nico, (1986), *The Emotions*, Cambridge University Press.

Fumerton, Richard, (2001), "Classical Foundationalism", *Resurrecting Old-Fashioned Foundationalism*, ed. Michael DePaul, Lanham, MD: Rowman & Littlefield.

Gallois, Andrè, (1996), *The World Without, the Mind Within*, Cambridge University Press.

Gazzaniga, M. S., and Smylie, C. S., (1984), "What Does Language Do for the Right Hemisphere?", ed. M.S. Gazzaniga, *Handbook of Cognitive Neuroscience*, New York: Plenum Press.

Gert, Joshua, (2006), "A Realistic Color Realism", *Australasian Journal of Philosophy* 84: 565-589.

Gertler, Brie, (2012), "Renewed Acquaintance", *Introspection and Consciousness*, eds. Declan Smithies and Daniel Stoljar, Oxford University Press.

Giardino, Valeria, and Greenberg, Gabriel, (2015), "Varieties of Iconicity", *Review of Philosophy and Psychology* 6 (1): 1-25.

Gibbard, Allan, (2003), *Thinking How to Live*, Harvard University Press.

Gibbard, Allan, (1990), *Wise Choices, Apt Feelings*, Harvard University Press.

Gigerenzer, Gerd, and Reinhard, Selten, eds., (2002), *Bounded Rationality: The Adaptive Toolbox*, MIT Press.

Gilbert, Margaret, (1989), *On Social Facts*, New York: Routledge.

Goodman, Nelson, (1968), *Languages of Art*, Indianapolis: Bobbs-Merrill.

Goodman, Nelson, (1967), "Seven Strictures on Similarity", *Experience and Theory*, eds. L. Foster and J. W. Swanson, University of Massachusetts Press.

Green, O. H., (1992), *The Emotions*, Dordrecht: Kluwer.

Greenberg, Gabriel, (2013), "Beyond Resemblance", *Philosophical Review* 122 (2): 215-287.

Greene, Joshua, (2008), "The Secret Joke of Kant's Soul", *Moral Psychology*, Vol. 3, ed. Walter Sinnott-Armstrong, MIT Press.

Greene, Joshua, (2003), "From Neural 'Is' to Moral 'Ought': What Are the Moral Implications of Neuroscientific Moral Psychology?", *Nature Reviews Neuroscience* 4: 847-50.

Greenspan, Patricia, (1988), *Emotions and Reasons*, New York: Routledge.

Griffiths, Paul, (1997), *What Emotions Really Are*, University of Chicago Press.

Grimm, Stephen, (2006), "Is Understanding a Species of Knowledge?", *The British Journal for the Philosophy of Science* 57: 515-36.

Guralnik, O., Schmeidler, J., and Simeon, D., (2000), "Feeling Unreal: Cognitive Processes in Depersonalization", *American Journal of Psychiatry* 157 (1): 103-109.

Hacker, P. M. S., (1987), *Appearance and Reality*, Oxford: Blackwell Publisher.

Haidt, Jonathan, (2012), *The Righteous Mind*, New York: Random House, Inc.

Haidt, Jonathan, (2001), "The Emotional Dog and Its Rational Tail", *Psychological Review* 108 (4): 814-834.

Harman, Gilbert, (1990), "The Intrinsic Quality of Experience", *Philosophical Perspectives 4*, *Action Theory and Philosophy of Mind*, ed. James Tomberlin, Atascadero: Ridgeview.

Harman, Gilbert, (1987), "(Nonsolipsistic) Conceptual Role Semantics", *Notre Dame Journal of Formal Logic,* ed. Ernest LePore, Academic Press.

Harman, Gilbert, (1982), "Conceptual Role Semantics", *Notre Dame Journal of Formal Logic* 28: 242-56

Harman, Gilbert, (1977), *The Nature of Morality: An Introduction to Ethics*, Oxford University Press.

Harman, Gilbert, (1973), *Thought*, Princeton University Press.

Harper, D. G. C., (1982), "Competitive Foraging in Mallards: Ideal Free Ducks", *Animal Behaviour* 30: 575-84.

Hawthorne, John, (2006), "Plenitude, Convention, and Ontology", *Metaphysical Essays*, Oxford University Press.

Helm, Bennett, (2001), *Emotional Reason: Deliberation, Motivation, and the Nature of Value*, Cambridge University Press.

Hilbert, David, (1987), *Color and Color Perception: A Study in Anthropocentric Realism*, Stanford, CA: Center for the Study of Language and Information.

Hill, Christopher S., (2009), *Consciousness*, Cambridge University Press.

Hinton, J. M., (1973), *Experiences*, Oxford: Clarendon Press.

Hopkins, Robert, (1998), *Picture, Image and Experience*, Cambridge University Press.

Horgan, Terence, (1984), "Jackson on Physical Information and Qualia", *Philosophical Quarterly* 32: 127-136.

Horgan, Terence, and Tienson, John, (2002), "The Intentionality of Phenomenology and the Phenomenology of Intentionality", *Philosophy of Mind: Classical and Contemporary Readings*, ed. David Chalmers, Oxford University Press.

Huebner, B., Dwyer, S., and Hauser, M., (2008), "The Role of Emotion in Moral Psychology", *Trends in Cognitive Science* 13 (1): 1-6.

Huemer, Michael, (2008), "Revisionary Intuitionism", *Social Philosophy and Policy* 25 (1): 368-392.

Huemer, Michael, (2005), *Ethical Intuitionism*, New York: Palgrave Macmillan.

Huemer, Michael, (2001), *Skepticism and the Veil of Perception*, Lanham, MD: Rowman & Littlefield.

Hume, David, (1742/2006), "Of the Standard of Taste", *Essays: Moral, Political, and Literary*, New York: Cosimo, Inc.

Hume, David, (1738-40/1975), *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd ed. revised by P. H. Nidditch, Oxford: Clarendon Press.

Hunter, E. C. M., Sierra, M., David, A. S., (2004), "The Epidemiology of Depersonalization and Derealisation: A Systematic Review", *Social Psychiatry and Psychiatric Epidemiology* 39: 9-18.

Hutcheson, Francis, (1726/2004), *An Inquiry into the Original of Our Ideas of Beauty and Virtue*, ed. W. Leidhold, Indianapolis: Liberty Fund.

Hyman, John, (2006), *The Objective Eye*, University of Chicago Press.

Inman, Ross, (2014), "Neo-Aristotelian Plenitude", *Philosophical Studies* 168 (3): 583-597.

Jackson, Frank, (2007), "Colour for Representationalists", *Erkenntnis* 66 (1-2): 169-85.

Jackson, Frank, (1998), *From Metaphysics to Ethics*, Oxford University Press.

Jackson, Frank, (1982), "Epiphenomenal Qualia", *Philosophical Quarterly* 32: 127-136.

James, William, (1890), *Principles of Psychology,* Vol. 2, Harvard University Press.

James, William, (1884), "What Is an Emotion?", *Mind* 9: 188-205.

Johnston, Mark, (2004), "The Obscure Object of Hallucination," *Philosophical Studies* 120 (1-3): 113-183.

Johnston, Mark, (2001), "The Authority of Affect", *Philosophy and Phenomenological Research* 63: 181-214.

Johnston, Mark, (1989), "Dispositional Theories of Value", *Proceedings of the Aristotelian Society* 63: 113-37.

Kahneman, Daniel, (2011), *Thinking Fast and Slow*, New York: Farrar, Straus and Giroux.

Kant, Immanuel, (1996), *Practical Philosophy*, trans. and ed. Mary J. Gregor, Cambridge University Press.

Keightley, David, (1996), "Art, Ancestors, and the Origins of Writing in China", *Representations* 56: 68-95.

Kelly, Thomas, (2005), "Moorean Facts and Belief Revision, or Can the Skeptic Win?", *Philosophical Perspectives* 19 (1): 179-209.

Kennett, Jeanette, (2006), "Do Psychopaths Really Threaten Moral Rationalism?", *Philosophical Explorations* 9 (1): 69-82.

Kenny, Anthony, (1963), *Emotion, Action, and Will*, London: Routledge & Kegan Paul Limited.

King, Peter, (2002), "Aquinas on the Passions", *Thomas Aquinas: Contemporary Philosophical Perspectives*, ed. Brian Davies, Oxford University Press.

Kitcher, Phillip, (1981), "Explanatory Unification", *Philosophy of Science* 48 (4): 507-531.

Kosslyn, S., Thompson, W., and Giorgio, G., (2006), *The Case for Mental Imagery*, Oxford University Press.

Kriegel, Uriah, (2013), "The Phenomenal Intentionality Research Program", *Phenomenal Intentionality*, ed. Uriah Kriegel, Oxford University Press.

Kriegel, Uriah, (2007), "Intentional Inexistence and Phenomenal Intentionality", *Philosophical Perspectives* 21 (1): 307-340.

Kuhn, Thomas, (1962), *The Structure of Scientific Revolutions*, University of Chicago Press.

Kulvicki, John, (2006), *On Images: Their Structure and Content*, Oxford University Press.

Kuppens, P., Tuerlinckx, F., Russell, J., and Barrett, L. F., (2013), "The Relation between Valence and Arousal in Subjective Experience", *Psychological Bulletin* 139 (4): 917-40.

Kvanvig, Jonathan, (2003), *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press.

Lange, C. G., (1885), "Om Sindsbevaegelser: et Psyko-fysiologisk Studie", København: Jacob Lunds.

Lazarus, Richard, (1991), *Emotion and Adaptation*, Oxford University Press.

LeDoux, Joseph, (1996), *The Emotional Brain*, New York: Simon & Schuster.

Lemos, Noah, (2004), *Common Sense: A Contemporary Defense*, Cambridge University Press.

Leslie, Sarah-Jane, (2011), "Essence, Plenitude, and Paradox", *Philosophical Perspectives* 25 (1): 277-296.

Levenson, Robert W., (2003), "Autonomic Specificity and Emotion", *Handbook of Affective Sciences*, eds. R. J. Davidson, K. R. Scherer, and H. H. Goldsmith, Oxford University Press.

Levine, Joseph, (2007), "Phenomenal Concepts and the Materialist Constraint," *Phenomenal Concepts and Phenomenal Knowledge*, eds. Torin Alter and Sven Walter, Oxford University Press.

Lewis, C. S., (1964/2012), *The Discarded Image*, Cambridge University Press.

Lewis, David, (1989), "Dispositional Theories of Value," *Proceedings of the Aristotelian Society* 63: 113–37.

Lewis, David, (1969), *Convention*, Harvard University Press.

Livy, (1912), *History of Rome*, trans. Rev. Canon Roberts, New York: E. P. Dutton and Co.

Loar, Brian, (2003), "Phenomenal Intentionality as the Basis of Mental Content", *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, eds. Martin Hahn and B. Ramberg, MIT Press.

Loar, Brian, (1990), "Phenomenal States", *Philosophical Perspectives* 4: 81-108.

Lopes, Dominic, (2005), *Sight and Sensibility: Evaluating Pictures*, Clarendon Press.

Lopes, Dominic, (1996), *Understanding Pictures*, Oxford University Press.

Lovejoy, Arthur O., (1936), *The Great Chain of Being*, Harvard University Press.

Lycan, William, (2007), "Moore's Antiskeptical Strategies", *Themes From G. E. Moore: New Essays in Epistemology and Ethics*, eds. Susana Nuccetelli and Gary Seay, Oxford: Clarendon Press.

Lycan, William, (2006), "The Meaning of 'Water': An Unsolved Problem", *Philosophical Issues* 16 (1): 184-199.

Lycan, William, (2001), "Moore Against the New Skeptics", *Philosophical Studies* 103 (1): 35-53.

Lycan, William, (1996), *Consciousness and Experience*, MIT Press.

Lyons, William E., (1980), *Emotion*, Cambridge University Press.

Mackie, J. L., (1977), *Ethics: Inventing Right and Wrong*, Pelican Books.

Marshall, G. D., (1976), "The Affective Consequences of 'Inadequately Explained' Physiological Arousal", unpublished doctoral dissertation, Stanford University.

Marshall, G. D., and Zimbardo, P. G., (1979), "Affective Consequences of Inadequately Explained Physiological Arousal", *Journal of Personality and Social Psychology* 37: 970-988.

Maslach, C., (1979), "Negative Emotional Biasing of Unexplained Arousal", *Journal of Personality and Social Psychology* 37: 953-969.

Matthen, Mohan, (1988), "Biological Function and Perceptual Content", *The Journal of Philosophy* 95: 5-27.

Matthews, Robert, (2007), *The Measure of Mind*, Oxford University Press.

McBrayer, Justin, (2010), "A Limited Defense of Moral Perception", *Philosophical Studies* 149: 305-320.

McDowell, John, (2008), "The Disjunctive Conception of Experience as Material for a Transcendental Argument", *Disjunctivism*, eds. A. Haddock and F. Macpherson, Oxford University Press.

McDowell, John, (1998a), "Aesthetic Value, Objectivity, and the Fabric of the World", *Mind, Value, and Reality*, Harvard University Press.

McDowell, John, (1998b), "Projection and Truth in Ethics", *Mind, Value, and Reality*, Harvard University Press.

McDowell, John, (1985), "Value and Secondary Qualities", *Morality and Objectivity*, ed. Ted Honderich, Routledge & Kegan Paul.

McDowell, John, (1979), "Virtue and Reason", *The Monist* 62 (3): 331-350.

McGinn, Colin, (1996), "Another Look at Color", *The Journal of Philosophy* 93: 537-555.

McLaughlin, B. P., (2009), "Monothematic Delusions and Existential Feelings", eds. T. Bayne and J. Fernandez, *Delusion and Self-Deception*, New York: Psychology Press.

Mendelovici, Angela, (2014), "Pure Intentionalism About Moods and Emotions", ed. Uriah Kriegel, *Current Controversies in Philosophy of Mind*, New York: Routledge.

Milona, Michael, (2016), "Taking the Perceptual Analogy Seriously", *Ethical Theory and Moral Practice* 19 (4): 897-915.

Montague, Michelle, (2009), "The Logic, Intentionality, and Phenomenology of Emotion", *Philosophical Studies* 145: 171-92.

Moore, G. E., (2013), *Philosophical Chapters*, New York: Routledge.

Moore, G. E., (1903), *Principia Ethica*, Cambridge University Press.

Mulligan, Kevin, (1998), "From Appropriate Emotions to Values", *The Monist* 81: 161-88.

Nagel, Thomas, (1986), *The View from Nowhere*, Oxford University Press.

Neu, Jerome, (1977), *Emotion, Thought, and Therapy*, Berkeley and Los Angeles: University of California Press.

Nisbett, Richard, and Ross, Lee, (1980), *Human Inference: Strategies and Shortcomings in Social Judgment*, Englewood Cliffs, NJ: Prentice-Hall.

Nozick, Robert, (1989), *The Examined Life*, New York: Simon and Schuster.

Nussbaum, Martha, (2001), *Upheavals of Thought: The Intelligence of Emotion*, Cambridge University Press.

Oddie, Graham, (2005), *Value, Reality, and Desire*, Oxford University Press.

Parfit, Derek, (2011), *On What Matters*, Vol. 2, Oxford University Press.

Peacocke, Christopher, (1987), "Depiction", *Philosophical Review* 96 (3): 383-410.

Pelser, Adam, (2014), "Emotion, Evaluative Perception, and Epistemic Justification," *Emotion and Value*, eds. Sabine Roeser and Cain Todd, Oxford University Press.

Pettit, Philip, (1991), "Realism and Response-Dependence", *Mind* 100 (4): 587-626.

Plantinga, Alvin, (1967), *God and Other Minds*, Cornell University Press.

Prinz, Jesse, (2008), *The Emotional Construction of Morals*, Oxford University Press.

Prinz, Jesse, (2004), *Gut Reactions: A Perceptual Theory of Emotions*, Oxford University Press.

Pritchard, Duncan, (2009), "Knowledge, Understanding, and Epistemic Value", *Epistemology (Royal Institute of Philosophy Lectures)*, ed. Anthony O'Hear, Cambridge University Press.

Pryor, James, (2000), "The Skeptic and the Dogmatist", *Noûs* 34: 517-549.

Pylyshyn, Zenon, (2003), *Seeing and Visualizing: It's Not What You Think*, MIT Press.

Pylyshyn, Zenon, (1981), "The Imagery Debate: Analogue Media Versus Tacit Knowledge", *Psychological Review* (88): 16-45.

Pylyshyn, Zenon, (1973), "What the Mind's Eye Tells the Mind's Brain: A Critique of Mental Imagery", *Psychological Bulletin* (80): 1-25.

Quine, Willard V., (1969), "Natural Kinds", *Ontological Relativity and Other Essays*, eds. Jaegwon Kim and Ernest Sosa, Columbia University Press.

Ratcliffe, Matthew, (2008), *Feelings of Being: Phenomenology, Psychiatry and the Sense of Reality*, Oxford University Press.

Reid, Thomas, (1872), *The Works of Thomas Reid*, Vol. 1, Edinburgh: Maclachlan and Stewart; London: Longman, Green, Longman, Roberts, and Green.

Reisenzein, R., (1983), "The Schachter Theory of Emotion: Two Decades Later", *Psychological Bulletin* 94 (2): 239-264.

Roberts, Robert C., (2013), *Emotions in the Moral Life*, Cambridge University Press.

Roberts, Robert C., (2003), *Emotions*, Cambridge University Press.

Roberts, Robert C., (1988), "What An Emotion Is: A Sketch", *Philosophical Review* 97: 183-209.

Robinson, Jenefer, (1979), "Some Remarks on Goodman's Language Theory of Pictures", *British Journal of Aesthetics* 19 (1): 63-75.

Ross, W. D., (1930), *The Right and the Good*, Oxford University Press.

Russell, Bertrand, (1914), "On the Nature of Acquaintance", *Monist* 24: 161-187.

Russell, Bertrand, (1912), *The Problems of Philosophy*, Oxford University Press.

Russell, Bertrand, (1910–11), "Knowledge by Acquaintance and Knowledge by Description", *Proceedings of the Aristotelian Society* 11: 108-128.

Russell, James, and Barrett, Lisa Feldman, (1999), "Core Affect, Prototypical Emotional Episodes, and Other Things Called Emotion: Dissecting the Elephant", *Journal of Personality and Social Psychology* 76 (5): 805-19.

Russell, James, (1980), "A Circumplex Model of Affect," *Journal of Personality and Social Psychology* 39: 1161-78.

Ryle, Gilbert, (1949), *The Concept of Mind*, Hutchinsons University Library.

Sacks, Oliver, (2012), *Hallucinations*, New York: Vintage Books.

Sacks, Oliver, (1995), *An Anthropologist on Mars*, New York: Vintage Books.

Scanlon, Thomas, (2014), *Being Realistic about Reasons*, Oxford University Press.

Scarantino, Andrea, (2014), "The Motivational Theory of Emotion", *Moral Psychology and Human Agency*, eds. Daniel Jacobson and Justin D'Arms, Oxford University Press.

Schachter, S., and Singer, J. K., (1962), "Cognitive, Social, and Physiological Determinants of Emotional State", *Psychological Review* 69: 379-399.

Schafer, Karl, (2013), "Perception and the Rational Force of Desire", *Journal of Philosophy* 110 (5): 258-281.

Schelling, Thomas, (1960), *The Strategy of Conflict*, Harvard University Press.

Schnall, S., Haidt, J., Clore, G., and Jordan, A., (2008), "Disgust as Embodied Moral Judgment", *Personality and Social Psychology Bulletin* 34: 1096-1109.

Schramme, Thomas, ed., (2014), *Being Amoral: Psychopathy and Moral Incapacity*, MIT Press.

Schwartz, N., and Clore, G. L., (1983), "Mood, Misattribution, and Judgments of Well-Being: Informative and Directive Functions of Affective States", *Journal of Personality and Social Psychology* 45: 513-523.

Searle, John, (1983), *Intentionality*, Cambridge University Press.

Selden, Daniel, (2013), *Hieroglyphic Egyptian*, University of California Press.

Shafer-Landau, Russ, (2003), *Moral Realism: A Defence*, Oxford University Press.

Siegel, Susanna, (2017), *The Rationality of Perception*, Oxford University Press.

Siegel, Susanna, (2011), *The Contents of Visual Experience*, Oxford University Press.

Sierra, Mauricio, (2009), *Depersonalization: A New Look at a Neglected Syndrome*, Cambridge University Press.

Sierra, M., Loperta, F., Lambert, F. V., Phillips, M. L., and David, A. S., (2002), "Separating Depersonalization and Derealisation: the Relevance of the 'Lesion Method'", *Journal of Neurology, Neurosurgery, & Psychiatry* 72: 530-532.

Simeon, Daphne, (2004), "Depersonalization Disorder: A Contemporary Overview", *CNS Drugs* 18 (6): 343-354.

Simeon, Daphne, and Abugel, Jeffrey, (2006), *Feeling Unreal: Depersonalization Disorder and the Loss of Self*, Oxford University Press.

Simeon, D., S. Gross, O. Guralnik, D. J., Stein, J. Schmeidler, and E. Hollander, (1997), "Feeling Unreal: 30 Cases of DSM-III-R Depersonalization Disorder", *American Journal of Psychiatry* 154 (8): 1107-1113.

Simeon, D., Guralnik, O., Hazlett, E. A., Spiegel-Cohen, J., Hollander, E., and Buchsbaum, M. S., (2000), "Feeling Unreal: A PET Study of Depersonalization Disorder", *American Journal of Psychiatry* 157 (11): 1782-1788.

Simon, Herbert, (1957), *Models of Man*, Oxford: Wiley.

Singer, Peter, (2005), "Ethics and Intuitions", *The Journal of Ethics* 9: 331-352.

Sinnott-Armstrong, Walter, (1991), "Moral Experience and Justification", *The Southern Journal of Philosophy* 29 (1): 89-96.

Sizer, Laura, (2006), "What Feelings Can't Do", *Mind and Language* 21 (1): 108-135.

Skyrms, Brian, (1996), *Evolution of the Social Contract*, Cambridge University Press.

Slovic, Paul, (2007), "'If I Look at the Mass I Will Never Act': Psychic Numbing and Genocide", *Judgment and Decision Making* 2 (2): 79-95.

Snowdon, P. F., (1979), "Perception, Vision and Causation", *Proceedings of the Aristotelian Society* 81: 175-92.

Sober, Elliot, (2015), *Ockham's Razors: A User's Manual*, Cambridge University Press.

Sober, Elliot, (2001), "What is the Problem of Simplicity?", *Simplicity, Inference and Modeling*, eds. A. Zellner, H. Keuzenkamp, and M. McAleer, Cambridge University Press.

Sober, Elliot, (1994), *From a Biological Point of View*, Cambridge University Press.

Solomon, Robert C., (1976), *The Passions*, Garden City, NY: Anchor Books.

Stampe, Dennis, (1987), "The Authority of Desire", *Philosophical Review* 96: 335–81.

Stampe, Dennis, (1986), "Defining Desire", *The Ways of Desire*, ed. Joel Marks, Chicago: Precedent.

Stich, Stephen, (1996), *Deconstructing the Mind*, Oxford University Press.

Stich, Stephen, and Warfield, Ted, eds., (1994), *Mental Representation: A Reader*, Blackwell.

Stocker, Michael, (1983), "Psychic Feelings: Their Importance and Irreducibility", *Australian Journal of Philosophy* 61: 5-26.

Sugden, Robert, (2011), "Salience, Inductive Reasoning, and the Emergence of Conventions", *Journal of Economic Behavior and Organization* 79: 35-47.

Sugden, Robert, (1986/2004), *The Economics of Rights, Co-operation, and Welfare*, 2nd ed., New York: Palgrave Macmillan.

Tappolet, Christine, (2012), "Emotions, Perceptions, and Emotional Illusions", *Perceptual Illusions: Philosophical and Psychological Essays*, ed. Clotilde Calabi, Palgrave-Macmillan.

Tolhurst, William, (1991), "On the Epistemic Value of Moral Experience", *Southern Journal of Philosophy* 29 (1): 67-87.

Tollefsen, Christopher, (2000), "McDowell's Moral Realism and the Secondary Quality Analogy", *Disputatio*: 1-13.

Tye, Michael, (2009), *Consciousness Revisited*, MIT Press.

Tye, Michael, (2008), "The Experience of Emotion: An Intentionalist Theory", *Revue Internationale de Philosophie* 62: 25-50.

Tye, Michael, (2000), *Consciousness, Color, and Content*, MIT Press.

Tye, Michael, (1995), *Ten Problems of Consciousness*, MIT Press.

Tye, Michael, (1986), "The Subjective Qualities of Experience", *Mind* 95 (377): 1-17.

Vance, Jona, (forthcoming), "Phenomenal Commitments: A Puzzle for Experiential Theories of Emotion", *The Ontology of Emotions*, eds. F. Teroni and H. Naar, Cambridge University Press.

Vargas, Manuel, and Nichols, Shaun, (2007), "Psychopaths and Moral Knowledge", *Philosophy, Psychiatry, and Psychology* 14 (2): 157-162.

Watkins, Michael, (2002), *Rediscovering Colors: A Study in Pollyanna Realism*, Dordrecht: Kluwer.

Wedgwood, Ralph, (2007), *The Nature of Normativity*, Oxford University Press.

Wiggins, David, (1987), "A Sensible Subjectivism?", *Needs, Values, and Truth*, Oxford University Press.

Wiggins, David, (1976), "Truth, Invention, and the Meaning of Life", *Proceedings of the British Academy* 62: 332-378.

Wittgenstein, Ludwig, (1953), *Philosophical Investigations*, trans. and ed. G.E.M. Anscombe, Basil Blackwell Ltd.

Wollheim, Richard, (1970), "Nelson Goodman's Languages of Art", *Journal of Philosophy* 67 (16): 531-539.

Wright, Crispin, (2004), "Warrant for Nothing (and Foundations for Free)?", *Aristotelian Society Supplementary Volume* 78 (1): 167-212.

Wright, Crispin, (1993), *Truth and Objectivity*, Harvard University Press.

Wright, Crispin, (1991), "Scepticism and Dreaming: Imploding the Demon", *Mind* 100 (1): 87-116.

Wright, Crispin, (1988), "Moral Values, Projection, and Secondary Qualities", *Aristotelian Society Supplementary* 63 (1): 1-26.

Wright, Crispin, (1987), "Realism, Antirealism, Irrealism, Quasi-Realism", *Midwest Studies in Philosophy* 12 (1): 25-49.

Yetter-Chappell, Helen, and Yetter-Chappell, Richard, (2013), "Mind-Body Meets Metaethics: A Moral Concept Strategy", *Philosophical Studies* 165 (3): 865-878.

Yong, Heming, Peng, Jing, and Tian, Bing, (2008), *Chinese Lexicography*, Oxford University Press.

Zagzebski, Linda, (2001), "Recovering Understanding", *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. Mattheus Steup, Oxford University Press.

Zajonc, Robert, (1980), "Feeling and Thinking: Preferences Need No Inferences", *American Psychologist* 35: 151-75.