

**Emotional Assessment and Emotion Regulation:
A Philosophical Approach**

by

Shai Madjar

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in The University of Michigan
2018

Doctoral Committee:

Professor Daniel Jacobson, Chair
Professor Sarah Buss
Professor Peter Railton
Professor Chandra S. Sripada

Shai Madjar

madshai@umich.edu

ORCID iD: 0000-0002-0158-9043

© Shai Madjar 2018

DEDICATION

This project is dedicated to my family.

ACKNOWLEDGMENTS

I am incredibly thankful to all of the people who have supported me, encouraged me, and challenged me to bring this dissertation to its current form.

I would like to thank each of my advisors (I secretly think of them as the “dream team”). Daniel Jacobson was an excellent chair, and I was very fortunate to work with him for a number of reasons. Dan was easy to work with, and I appreciated his direct approach. He always seemed to have a great sense of just what I needed to do to turn an initial draft of my ideas into a successful chapter – and he was always right about this. Dan also helped me to think as carefully as I could about the nature of fittingness, and without his expertise on this subject, I don’t think that I would have developed my ideas as far as I did. I also greatly benefitted from Sarah Buss’s sharp eyes. Sarah helped me to recognize some of the most important issues with my arguments and consistently helped me to find more nuanced and insightful positions. I am thankful to have had her support for some time now, both professionally and academically (I also had the great fortune of working with her on my undergraduate thesis). Chandra Sripada has also been supporting me for some time, ever since I began to consider pursuing the path of an MD-PhD. Chandra also helped me to think through my very initial ideas for this project. And he has continued to advise me throughout this process. Chandra has been great in helping me to accentuate, in each chapter of this dissertation, the key moves that I am trying to make. And finally, I am very thankful that Peter Railton is a part of this project. Peter has been a great influence on my thinking, and it was he that initially inspired me, as an undergraduate, to think seriously about emotions by drawing on perspectives from both philosophy and psychology. Peter, too, has been advising me for some time – he guided me through my undergraduate thesis project, and his feedback on this project has been very valuable, helping to shape my thinking on a number of key issues that I needed to clarify.

I am also extremely thankful to my family. My family has supported me every step of the way. No matter what my own anxiety about the project had to say, their faith in me was unwavering.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGMENTS	iii
ABSTRACT	v
Introduction	1
Chapter I. Emotions, Action-Responses, and the Nature of Fittingness	5
Chapter II. Emotion Regulation and Evaluative Understanding	38
Chapter III. On the Distinctive Final Value of Negative Emotions	78
Conclusion	112
Bibliography	118

ABSTRACT

This dissertation contains three standalone chapters, each of which addresses a different philosophical issue related to emotional assessment or emotion regulation. But each of these chapters contributes to the larger goal of understanding when and how we should regulate our emotions.

In chapter 1, I examine what it means to say that an emotion is fitting. I argue that in order for an emotion to be fitting, it must do more than correctly represent its object; it must also mobilize the individual to correctly respond to this object. My analysis, which also leads me to argue that action-responses can be fitting, suggests that fittingness is not fundamentally about the correctness of mental representations.

In chapter 2, I investigate the relationship between emotions, emotion regulation, and evaluative understanding. I argue that if our goal is evaluative understanding, we should regulate our emotions in a particular, organized way, that involves both engagement with our emotional concerns, as well as disengagement.

In chapter 3, I investigate the claim that fitting negative emotions have a distinctive final value, for epistemic or moral reasons, that calmer mental states cannot possess. I argue that this claim is false. We can, and often do, downregulate our emotions without losing anything of distinctive epistemic or moral final value.

I conclude the dissertation by briefly describing the structure of a practical, normative model for emotion regulation.

Introduction

Suppose that you are anxious about some future threat, sad about some loss or setback, or angry about some perceived injustice. What should you do while in the grip of this emotion? Should you allow it to guide your thoughts and actions? Or should you regulate this emotion? But if you do choose to regulate your emotion, how should you do so? What sort of emotion regulation techniques should you rely upon? In order to properly answer such questions, one must address a number of philosophical issues about emotional assessment and emotion regulation. You might worry, for instance, about what you could lose in regulating your emotion: a fitting response, a response that might promote your evaluative understanding, a response that, although painful, may help you to *feel* the importance of some concern, or to express how much you care about it. You might also wonder whether there are certain forms of emotion regulation that are, in light of such worries, more epistemically or morally responsible. In my dissertation, I examine these issues in order to clarify the value and wisdom of emotion regulation, in its various forms.

In Chapter 1, I investigate the nature of fittingness. When we endorse an emotion as fitting, what is the nature of this endorsement? I argue against the standard view in the philosophy of emotion, according to which an emotion is fitting if and only if it correctly represents its target - call this the *recognitionist view* of emotional fittingness. This view fits in nicely with a more general ambition to understand the fittingness of a response in terms of a

correct mental representation. However, I consider two problem cases that lead me to reject this type of view. First, I argue that in order to be fitting, emotions must do more than correctly represent their target values. In order to be fitting, emotions must also correctly mobilize us to respond to these values. Second, I argue that, perhaps surprisingly, even action-responses can be assessed for fittingness. Just like emotions, beliefs, and desires, action-responses can be supported by the *wrong kind of reason*. But this suggests that the fittingness of a response is not essentially about the correctness of mental representations. Instead, fittingness is a distinctively narrow form of assessment that can be applied to any response. I suggest that we can understand fittingness either as a normative primitive, or in terms of reasons: a fitting response is a response that is supported by the right kind of reasons.

In Chapter 2, I investigate the relationship between emotions, emotion regulation, and evaluative understanding. Emotions can enhance our evaluative understanding by mobilizing directed reflection: by worrying about some threat, ruminating about some loss, or simmering about some injustice, we can enhance our understanding of the threat, loss, or injustice in question. But notoriously, emotional reflection can also lead us astray. If our goal is evaluative understanding, then, we must make room for emotion regulation. But which forms of emotion regulation should we rely upon, if our goal is evaluative understanding? In this chapter, I distinguish between engaged forms of emotion regulation, which keep us engaged with our emotional concern (e.g. certain forms of reappraisal), and disengaging forms of emotion regulation, which regulate emotional experience by leading us to direct attention away from the emotional concern in question (e.g. many forms of meditation). I consider, but then reject, the *engagement view*, according to which engaged forms of emotion regulation

characteristically enhance evaluative understanding, whereas disengaging forms of emotion regulation hinder or detract from evaluative understanding. Against this view, I argue that disengaging forms of emotion regulation can play a vital role in enhancing evaluative understanding. I propose a practical model that can help us to decide when to rely on engaging forms of emotion regulation and when to rely on disengaging forms of emotion regulation, if our goal is evaluative understanding.

In Chapter 3, I investigate the final value of painful negative emotions. A number of philosophers argue that painful negative emotions, when fitting, possess a distinctive final value, for epistemic or moral reasons, that calmer mental states cannot possess. For example, it is argued that only by being angry at injustice, only by grieving over significant losses, and only by feeling appropriately guilty about personal wrongdoing can we fully appreciate the relevant concerns (injustice, loss, and personal wrongdoing), or fully demonstrate that we care about them. Call this the *distinctive final value* thesis (DFV). In this chapter, I argue that DFV is false, though I also explain why we might nevertheless find it difficult to resist. Now, I do not deny that painful negative emotions, when fitting, possess final value for epistemic or moral reasons. But I argue that this value is not distinctive; calmer mental states can possess the very same final value. The outcome of this debate has important practical implications for emotion regulation. If DFV is true, then we always have at least a pro tanto reason not to regulate our painful, yet fitting negative emotions. If such reasons are at all weighty, then it may be that we ought to regulate our emotions far less often than we might have thought. By contrast, if DFV is false, then an important normative obstacle for emotion regulation is removed, and the way we think about our emotions may have to change.

I shall conclude the dissertation by briefly discussing the implications these chapters have for when and how we should regulate our emotions. I shall then briefly describe the structure of a practical, normative model for emotion regulation that is informed by these considerations. This model will emphasize the importance of emotion regulation for 1) enhancing our evaluative understanding and 2) helping us to act in accordance with our understanding.

Chapter I: Emotions, Action-Responses, and the Nature of Fittingness

Introduction

When we endorse a response as fitting, what is the nature of this endorsement? In order to investigate this question, I shall begin by examining the nature of emotional fittingness. Since emotions are complex responses, if we can understand fittingness in the case of emotions, this might help us to understand the nature of fittingness more generally.

According to a standard view in the philosophy of emotion, an emotion is fitting if and only if it correctly represents its object.¹ Fear, for instance, is fitting just in case it correctly represents an object as dangerous. Anger is fitting just in case it correctly represents a moral violation of some kind.² And so on – call this the *recognitional view* of emotional fittingness. The recognitional view of emotional fittingness is attractive not only for its simplicity, but also because it dovetails nicely with our conception of how other mental states, such as beliefs,

¹ I will discuss this view in more detail later in Section 1, and I will suggest that it is quite prevalent. In their earlier work, D'Arms and Jacobson (2000) seemed to endorse this view: "to call an emotion appropriate is to say that the emotion is fitting: it accurately presents its object as having certain evaluative features" (Ibid p. 65). When characterizing this type of view, philosophers seem to use the terms 'presents' and 'represents' interchangeably (for an exception, see Dokic and Lemaire, 2013).

² See, for instance, Srinivasan (2018, p. 128-130).

come to be fitting: namely, by correctly representing the world. We might, in other words, be tempted by a unified theory of fittingness, according to which a fitting response is a response that correctly represents some aspect of reality. For convenience, we might refer to this account as the *recognitional view*, without qualification.

In this paper, I shall argue that the recognitional view is false. But my arguments against the recognitional view will help to elucidate the true nature of fittingness assessments.

I shall begin by examining the recognitional view of emotional fittingness, which is quite prevalent in the literature.³ I shall argue that the recognitional view of emotional fittingness is mistaken, generating the first problem case for the more general version of the recognitional view. In particular, I shall argue that the fittingness conditions of emotions go beyond correct representation. In order to be a fitting response to its target value, an emotion must not only correctly represent the value in question, it must also mobilize the individual in a way that is correctly guided by the value. What does this require? It requires, as I shall explain, that the emotional response be fully supported by the reasons provided by the value, the *right kind of reasons* – as opposed to the reasons provided by extraneous considerations, the *wrong kind of reasons*.⁴

³ See section 1 for details. Now, the recognitional view of emotional fittingness relies on the premise that emotions present or represent the world as having certain evaluative features, an assumption that has recently been criticized in a number of different ways. See, for instance, Whiting (2012), Dokic and Lemaire (2013), Schroeter, Schroeter, and Jones (2015), and Deonna and Teroni (2015). However, in this paper, I will grant this assumption to the recognitionalist for the sake of argument. I will argue that even if emotions represent values, the correctness of these representations is not sufficient for the fittingness of the response as a whole.

⁴ For discussion of the wrong kind of reason problem, see D’Arms and Jacobson (2000) and Rabinowicz and Rønnow-Rasmussen (2004). See Gertken and Kiesewetter (2017) for a review of the literature.

Having considered the nature of emotional fittingness, I shall extend the discussion to another set of responses, *action-responses*. Perhaps surprisingly, I shall argue that action-responses can be assessed for fittingness. Just as fear is a fitting response to danger, I shall suggest that calm avoidance is a fitting response to danger. And just as anger is a fitting response to injustice, I shall suggest that calm resistance is fitting too. But the fittingness of action-responses does not seem to be based on the correctness of mental representations. This, then, is the second problem case for the general version of the recognitional view; if action-responses can be assessed for fittingness, then this suggests that fittingness is a phenomenon that is not essentially about mental representations at all. In developing this objection, I shall argue that when it comes to action-responses, fittingness assessments are familiar and intuitive, and they bear the hallmark of such attributions: just like emotions, action-responses can be supported by the *wrong kind of reason*.⁵

⁵ Mark Schroeder (2010) has argued that *some* actions can be supported by the wrong kind of reason, since they belong to practices that possess internal standards of correctness (for example, there are correct ways to make a chess move, tie a knot, and so on). Indeed, Schroeder suggests that reasons of the wrong kind are easy to come by, requiring only the presence of a relevant standard of correctness. I agree with Schroeder on this point, but I shall point out that standards of correctness are even easier to come by than Schroeder suggests. I shall argue, for instance, that *any* action-response can be assessed for fittingness, and that such assessments are not dependent on special practices with internal standards of correctness. This is because responses can always be assessed in relation to circumscribed *external* standards, determined by what a given response is a response *to*.

In a PEA Soup blog post, Rabinowicz and Ronnow-Rasmussen (2012) also raise the idea that actions can be supported by the wrong kind of reason. It is unclear to me just how general their account is meant to be, but my account shall differ from theirs in at least one important way. I make the distinction between actions and action-responses, and I claim that in contrast to the former, the latter can always be assessed for fittingness. Rabinowicz and Ronnow-Rasmussen concede that actions lack one of the typical earmarks of wrong-kind-of-reason susceptibility (namely, it is characteristically difficult to have an attitude for the wrong kind of reason. However, it is not difficult to perform an action for the wrong kind of reason). By contrast, I shall argue that action-responses do possess this typical earmark.

Heuer (2010) and Hieronymi (2005, 2013) argue that actions cannot be done for the wrong kind of reason. I shall agree that without special reference to a circumscribed standard, actions cannot be done for the wrong kind of reason, but I shall argue that any action-response can be supported by the wrong kind of reason. See Section 3.

Now, some philosophers argue that the “wrong kind of reasons” aren't really reasons at all (see, for instance, Gibbard, 1990, Skorupski, 2007, Parfit, 2011, and Way, 2012). This so-called *wrong kind of reason skepticism* is compatible with my view. I do not insist that wrong-kind-reasons are genuine reasons, only that there

Of course, this all means that the notion of fittingness requires further examination. But what sort of account could accommodate the fittingness of both emotions and action-responses? In this paper, I shall suggest two options. The first is to take fittingness as a normative primitive, as recently suggested by some philosophers.⁶ The second option, compatible with a more familiar reasons-first approach, would be to understand a fitting response simply as a response that is supported by the right kind of reasons.

Either approach would be compatible with what I take to be the true nature of fittingness. In my view, fittingness assessments are distinctively narrow in focus; to endorse a response as fitting is to endorse it in relation to a circumscribed standard of correctness, in a way that is insensitive to *extraneous* considerations, or what are often called the wrong kind of reasons. Part of the novelty of my account shall be the claim that such standards of correctness are ubiquitous. They are not limited to the internal standards of correctness that characterize certain mental states.⁷ For *any* response, insofar as it is a response to some external, circumscribed consideration, can be narrowly assessed as being a fitting or unfitting response to *that* particular consideration. I shall suggest that fittingness assessments are defined by this narrowness, and not by any specific focus on the correctness of mental representations.

is a distinction to be made between reasons that genuinely bear on the fittingness of an action-response, and either wrong-kind-of-reasons or *non-reasons* that do not. Still, for ease of exposition, I will continue to speak of the "wrong kind of reasons," with the understanding that these may not be genuine reasons.

⁶ See Chappell (2012), McHugh & Way (2016), Howard (forthcoming). See Rowland (2017) for a recent criticism.

⁷ Nor are they limited to internal standards generated by special practices, such as the rules of a game, the rules for setting a table, etc. – see Schroeder (2010).

Finally, an important point of clarification. Throughout the paper, I shall be regarding fittingness as a genuine normative phenomenon that is deserving of study, just as reasons and values are deserving of study. I reject the view that ‘fitting’ is merely a philosopher’s term of art, and I think that I have excellent reasons for doing so. 1) I shall be demonstrating, through examples, that we have strong intuitions about fittingness in cases that have not been explored in the literature, suggesting that fittingness is a phenomenon that we can all pre-theoretically recognize. Moreover, the notion of fittingness that is involved in these cases, as we shall see, appears to be the same one that has interested philosophers. 2) Philosophers themselves are already calling out for an account of fittingness, and a number of different theories have been suggested.⁸ 3) In particular, some philosophers have proposed that fittingness may turn out to be the most primitive normative notion, to which all others reduce.⁹ Far from being a term of art, ‘fittingness’ may pick out one of the most fundamental normative phenomena. For these reasons, I will be taking seriously the notion of fittingness in the investigation that follows.

In Section 1, I shall describe the recognitional view of emotional fittingness and suggest that it is quite prevalent. In section 2, I shall argue that emotions must do more to be fitting than correctly represent their target values, and so, that the recognitional view of emotional fittingness is false. In Section 3, I shall argue that action-responses can be assessed for fittingness, a second problem case for the recognitional view. In Section 4, I shall outline

⁸ See Howard (forthcoming) for an overview. I shall be considering the most important alternatives in this paper: the recognitional view (also known as the “alethic view”), the reasons-first view, and the fittingness-first view.

⁹ See footnote 6.

possibilities for an alternative conception of fittingness, elaborate on some important points, and consider difficult cases.

1. The Recognitional View of Emotional Fittingness

Let us begin our investigation into the nature of fittingness. The notion of *fit* invites an image of two puzzle pieces, *A* and *B*, whose structure is such that the two pieces can be joined together without difficulty or resistance. In the case of *emotional* fittingness, *A* is thought to be the value, or evaluative property, or significant aspect of a situation to which the emotion, *B*, is a response. Fear is a fitting response to danger. Anger is a fitting response to a moral violation of some kind. And so on. What is it that makes these two pieces fit together? A standard answer in the literature seems to be that emotions *represent* values or evaluative properties. And so, emotions are fitting when their representations match up with or *fit* the evaluative properties that are actually out there. Call this the *recognitional view* of emotional fittingness:

Recognitional View of Emotional Fittingness: An emotional response is fitting if and only if it correctly represents its object.

It is difficult to say just how widespread the recognitional view of fittingness is in the philosophy of emotion, since philosophers that rely on the notion of fittingness or appropriateness do not usually offer a precise characterization of the concept. Nevertheless, the view seems prevalent enough to suggest that it is the dominant view in the philosophy of emotion literature. For example, the recognitional view lends itself well to judgmentalist theories of emotion, on which emotions are thought to be judgments of a special kind. Judgments are typically assessed in terms of whether they are true or false, whether they

correctly represent some state of affairs. If emotions are judgments, then perhaps emotions too should be assessed in this way. For example, Nussbaum (2001, p. 47) writes:

Emotions can be true or false ... Often, trying to avoid the implication that statements about value can be true or false, we say instead that they are “appropriate” or “inappropriate.” The language of appropriateness, however, confuses the issue, since things can be appropriate or inappropriate in many different ways: it doesn’t single out the aspect of value-correctness that we want to single out.

Nussbaum is here clearly searching for a term like ‘fitting,’ a concept that picks out a narrow form of appropriateness, a correct relationship between a response and a value. For Nussbaum, an emotion is appropriate in this special sense (that is, fitting) when it is *true*, or “value-correct,” just as the recognitional view requires. Similarly, for Solomon (2007, p. 181), “the target of an emotion is getting its object right”¹⁰ – that is, correctly representing the evaluative object.

Now, if the recognitional view lends itself well to judgmentalist theories of emotion, it also lends itself well to perceptual theories of emotion, and for the same reason. According to perceptual theories, emotions are analogous to perceptions. Perceptions can be correct or incorrect, they can represent a situation accurately or inaccurately. And so, if emotions are analogous to perceptions, then perhaps emotions should also be assessed in this way, in accordance with the recognitional view. For example, Doring (2007, p. 384) writes: “the evaluation [an emotion] contains is a representation ... representational content aims at correctness (so to speak), which is at fitting the world,” and “This is in analogy to the content of

¹⁰ Solomon is quick to point out, however, that our emotions finding their “targets” is just one part of a “larger picture,” since we can also assess emotions in terms of their prudential value (Ibid).

a sense perception” (Ibid p. 377). Similarly, Tappolet (2005, p.229) writes that “an emotion like fear has correctness conditions that are similar to the correctness conditions of perceptual states. Fear is correct or appropriate in so far as its object is really dangerous.”¹¹

The recognitional view is also recommended by those who would challenge the analogy between perception and emotion. Deonna and Teroni, for instance, point out that we give reasons for our emotional responses, but not our perceptions. Nevertheless, it is suggested (2012, p. 64) that “the predicates ‘appropriate’ and ‘justified’ mean, when applied to the emotions, what they customarily mean in other areas of discourse, namely that there are good and undefeated reasons for representing the facts as these emotions do.”

The recognitional view of emotional fittingness is also endorsed in D’Arms and Jacobson (2000), an important paper in which fittingness is clearly distinguished from other forms of emotional assessment: “Emotions present things to us as having certain evaluative features. When we ask whether an emotion is fitting, in the sense relevant to whether an object is ϕ , we are asking about the correctness of these presentations” (Ibid, p. 72). I conclude that the recognitional view of fittingness is prevalent in the philosophy of emotion and is perhaps the dominant view of emotional fittingness. But is the recognitional view of emotional fittingness true? In the next section, I will argue that it is not.

¹¹ See also Tappolet (2011).

2. Emotional Responses

Emotions do more than represent their objects as having evaluative features. Emotions also mobilize individuals to respond to these evaluative features.¹² And just as the emotion's representation can be correct, the way in which it mobilizes the individual can also be correct. This is because the value in question *calls for* certain responses, and emotional responses can heed, or fail to heed, this call. As I shall now illustrate, this claim is borne out by our intuitions: fear seems to be a fitting response to danger not only because it represents dangers, but also because it mobilizes threat-avoidance. When this mobilization is sufficiently defective, the emotional response will not seem fitting.

Suppose, for instance, that Fahad encounters a bear in his backyard, and so is immediately consumed by fear. The danger posed by this bear calls for some kind of avoidance response. And indeed, Fahad's fear answers this call and endeavors to mobilize him to escape the bear. Not only does Fahad's fear represent the bear as dangerous, it also directs his cognition, motivation, and physiology in preparation for escape. For example, Fahad, slowly backing away, immediately thinks of his back door as a means of escape, but he also considers hiding behind a big bush, which is closer. His memory also kicks into gear, and he seems to remember that there's some kind of advice people give about how to deal with such bear

¹² Emotions are complex responses, involving changes in attention, thinking, feeling, memory, motivation, and physiology. Perhaps these various parts of the emotion are not all *necessary* parts. But intuitively, they are all still parts of the emotion. All four legs of a table, for instance, are parts of the table, even though no one of them is a necessary part. When deciding whether a table fits a room, we have to take into account all of its parts, not just the necessary parts (if there are any). The same is true for emotional fittingness. Now, some may wish to insist that the various changes associated with emotions are not even contingent parts of these emotions (only the representation of value is the emotion proper). As I note in just a moment, I find such a view unduly restrictive, but the larger point that I am trying to make here does not depend on my being right about this. If one wishes, one can speak of a *composite* of the emotion proper and its downstream changes. And this emotional composite, I argue, is assessable for fittingness, and in a way that depends on more than just a correct representation.

encounters, but he's forgotten exactly what it is. His heart rate and respiration elevate, preparing him for a quick dash (or a fight for survival), if necessary. Each of these response components seems correctly guided by danger, by this threat of harm. It's not that these responses are guaranteed to be useful – for example, Fahad's memory didn't actually help; and perhaps this is a very hungry bear, and there was never much chance of escape anyway. Still, each component of Fahad's fear answered the call for avoidance. The danger posed by the bear generated a *reason* to try to remember how to escape bears, a reason to search for escape routes, a reason to prepare physiologically for flight, and so on. Because Fahad's fear response, as a whole, responded to these reasons, his response as a whole was fitting.

By contrast, consider the case of Gad. Gad has a defective fear response that correctly represents dangers, but then mobilizes him in a way that has nothing to do with avoidance. In particular, whenever Gad perceives dangers, he does feel frightened, and he dearly wishes to get away, but the downstream cognitive changes associated with his emotion do not direct him towards plans for escape. Instead, he begins to think about vegetables, and he suddenly finds himself remembering various kinds. So when Gad sees a bear in his backyard, his fear correctly represents the bear as dangerous. But where Fahad's attention would, as a result, immediately turn to possible routes of escape, Gad is paralyzed, shaking in his boots, thinking about broccoli. And *that* does not seem to be a fitting response to danger. Although Gad's defective fear response correctly represents the bear as dangerous, his defective fear fails to heed the danger's call for avoidance.

Now, it might be objected that the downstream cognitive changes associated with an emotion are not part of the emotion proper. And so, Gad does feel fear, but this fear does not

itself contain the defective cognitive changes. But this objection requires a very narrow view of emotion that I think we should reject, for even if the cognitive, motivational, and physiological changes associated with some emotion are not necessary parts of the emotion, they are still contingent parts of the emotional response. Happily, however, the larger point that I am trying to make here does not depend on my being right about this. Even if there is a distinction to be made between an emotion proper and its downstream changes, it seems that we can still assess the *composite* of these two elements for fittingness. Call this the “composite emotional response,” if you like (but ultimately, I think this phrase has the same reference as “the emotional response”). Intuitively, Gad’s composite emotional response, which involves fearfully thinking about vegetables, is not a fitting response to danger, even if his emotion proper correctly represents the danger.

It might now be doubted that Gad’s emotion really did represent the bear as dangerous, given his strange, vegetable-related cognitions. But I do not think this doubt has any basis. Just as Fahad represents the bear as dangerous, and then fearfully thinks about escape, Gad represents the bear as dangerous, and then fearfully thinks about vegetables. After the fact, Gad might explain how he felt by saying, “I was so afraid of the bear, but I couldn’t help but think about vegetables!”¹³ In other words, I see no problem in stipulating that Gad’s fear represented the bear as dangerous, and then mobilized him inappropriately. And so, I suggest that Gad’s defective response was unfitting, even though it correctly represented the danger. If

¹³ And Gad would deny, if you’re worried about this, that his thoughts about vegetables were somehow about the bear as well. No, the bear seemed *dangerous*, not related to vegetables at all. But in response to this danger, Gad experienced intrusive thoughts about broccoli.

that's right, then we still need a conception of fittingness that goes beyond correct representation.

Or do we? It might be objected that the difference between Fahad's response and Gad's response is that Fahad's was prudentially valuable, and Gad's response was not. That's why it seems that Fahad's response was fitting, but Gad's was unfitting. But this is not the right explanation. As I noted in the example, Fahad's memory response didn't actually uncover anything, and perhaps the escape routes he identified would fail him in the end (the back door is locked, and the bush is a terrible hiding spot). Still, the threat posed by the bear provided Fahad with a *reason* to try to remember, a *reason* to search for escape routes, and so on. Fahad's response as a whole was supported by these threat-related reasons, and this is why his response was fitting. Gad's defective vegetable-response, by contrast, was not supported by the reasons provided by the threat. This is why his response was not a fitting response to the threat.

But why use the term 'fitting' here? Why not just say that Fahad's response was supported by reasons, whereas Gad's response was not? The answer is that not just any reason will do. In saying that Fahad's response was fitting, I mean to say that his response was supported by the *right kind* of reasons. To illustrate, suppose that an eccentric billionaire is willing to pay a large sum of money to anyone who, in the presence of a bear, thinks about or tries to remember facts about vegetables. Intuitively, both Fahad and Gad now have a reason to engage in this odd type of response. But even if that's true, that doesn't make this response a *fitting* response to danger, to the threat posed by the bear! Even though Gad's response is

now supported by a reason, it is still not a fitting response to danger. The billionaire's offer is what philosophers call the "wrong kind of reason," a reason for some response that does not bear on whether that response is *fitting*. Gad's response is a response to *danger*. If we want to know whether Gad's response was fitting, therefore, we have to assess his response in relation to the danger, which provides us with a *circumscribed standard* for assessment. Any consideration that is external to this standard will strike us as the wrong kind of reason. The billionaire's offer is a case in point. Now, it is notoriously difficult to articulate systematic criteria for distinguishing between the right and the wrong kinds of reasons.¹⁴ But intuitively, the threat posed by the bear, which is what Fahad and Gad were responding to, provided Fahad and Gad with the right kind of reasons. For instance, Fahad and Gad had a reason of the right kind to try to remember bear-related safety facts. This is why Fahad's memory response was fitting – it was supported by the right kind of reason.

I shall discuss this phenomenon further in the next section on action-responses. For now, it is important to note a crucial point. Emotional representations are not the only components of emotional responses that can be assessed in relation to a circumscribed standard of correctness. As a result, they are not the only components that can be assessed for fittingness. As we have seen in the cases of Fahad and Gad, it seems that the emotional response as a whole can be assessed for fittingness. This point, of course, applies not just to fear, but to other emotions too, such as anger, shame, and sadness. Each emotion mobilizes the individual cognitively, motivationally, and physiologically. In order to be fitting, the

¹⁴ See Rabinowicz and Rønnow-Rasmussen (2004) and Gertken and Kiesewetter (2017) for a review.

emotional response must not only correctly represent a value, it must also answer to the reasons provided by this value. The recognitional view, however, insists that a correct representation is sufficient for fittingness. For this reason, the recognitional view is mistaken.

3. Action-Responses

In the previous section, I argued that the fittingness conditions of emotional responses seem to go beyond correct representation. This is the first problem case for the recognitional view of fittingness. In this section, I shall argue that, like emotions, *action-responses* can be assessed for fittingness in a way that seems familiar and intuitive. This shall serve as a second problem case for the recognitional view, which cannot account for the fittingness of action-responses, given its focus on correct representation.

Let us begin with the case of the seasoned soldier Sasha, who no longer feels fear on the battlefield. Although Sasha no longer feels fear, her experience means that she is exquisitely sensitive to dangers. Noticing dangerous movements in the enemy troop's position, she alters the position of her own troop. It seems natural to say here that Sasha's response is fitting – a fitting response to danger. Now, this is not because we think that her response is guaranteed to be *effective*. It could well be that Sasha's chosen maneuver will not ultimately succeed. Still, her attempt to redirect her troop to safety seems to be a fitting response to danger (her response is correctly guided by danger's significance). Now, an advocate of the recognitional view might wish to capture the fittingness of Sasha's response by pointing to the fact that she correctly judged danger to be present. But Sasha's correct judgment cannot explain the fittingness of her response as a whole. After all, had Sasha correctly recognized danger, but then directed her

troops to recklessly run towards the threat, this would not have been a fitting response to danger.

Consider another case. Misa is a staunch advocate of equality. She routinely organizes marches, protests, and other events in an effort to combat various injustices. When Misa protests a specific injustice in this way – a particularly glaring instance of discrimination, say – this seems to be a fitting response to injustice. This is not because Misa’s response is guaranteed to be morally valuable, all things considered. For all we know, Misa’s specific approach might lead to more harm than good. Rather, Misa's efforts to resist the injustice seem fitting because injustices call for such resistance, such confrontation.¹⁵ Now, once again, the fittingness of Misa's response cannot be explained simply in terms of her correct recognition of injustice. Misa could correctly recognize the injustice, but then stay at home and do nothing. And that, intuitively, would not be a fitting response to injustice.

One last example. Suppose that Gloria has just received a poor grade on her exam. And let us grant for the purpose of the example that her score is quite shameful; after all, the test really wasn’t difficult, but Gloria had decided to “wing it” – she hadn’t studied at all. Aware of her classmates around her, Gloria folds her exam paper over, so that no one else can see her grade. This seems to be a fitting response to the shamefulness of her score, which reflects poorly on her, and which she has a reason to hide. Once again, it’s hard to see how the

¹⁵ Injustice has a different significance for those performing unjust acts, or for those who are about to perform an unjust act. Such individuals have reason to refrain from what they are doing, or to develop an intention to refrain. This is why it is fitting for an individual to feel guilty upon performing an unjust act. Guilt motivates the individual to refrain from performing such acts in the future. See Gibbard (1990) for a view on which the significance of wrongness is different for those performing wrong actions and those witnessing wrong actions – it is fitting for the former to feel guilt and for the latter to feel resentment.

recognition view could capture the fittingness of Gloria's response. Had Gloria recognized the shamefulness of her score, but then advertised it loudly to the class, this would not have been a fitting response to the score's shamefulness.

I hope that these three examples illustrate that action-responses can be assessed for fittingness in a way that is familiar and intuitive. But the fact that action-responses can be assessed in terms of fittingness can also be demonstrated more methodically, as we have seen: we can show that action-responses can be supported by the *wrong kind of reason*, a hallmark of fitting responses. A fitting response, after all, is a response that fits a circumscribed value or consideration, and this fit will be insensitive to extraneous considerations, or to the wrong kind of reasons. If we can identify such wrong-kinds-of-reasons for action-responses, then, we can show that action-responses possess this core feature of fittingness. And indeed this is possible.

You might have noticed that I have been careful to speak of *action-responses*, as opposed to *actions* simpliciter. There is a method to my madness. In general, an action, such as *running*, can be performed for all sorts of reasons, and any of these reasons can contribute to the practical justification of the action. The presence of danger, for instance, is a reason for running, but then so is the offer of an eccentric billionaire. Neither of these reasons seems to be a reason of the wrong kind.¹⁶ By contrast, an *action-response* is an action that is undertaken in response to some specific consideration – for example, *running in response to danger*.

Perhaps all actions, when actually performed by agents, are action-responses, since actions are

¹⁶ See Heuer (2010) and Hieronymi (2005, 2013) for the view that actions cannot be supported by the wrong kind of reason. But see Schroeder (2010) for the view that *some* actions are supportable by the wrong kind of reason, since they belong to practices with internal standards of correctness (e.g. executing an opening in chess). By contrast, I shall be arguing that *any* action-response can be assessed for fittingness.

always done for some reason. Nevertheless, if we are careful to assess action-*responses* as opposed to actions, we will find that they can be supported by the wrong kind of reason. This is because the consideration that the individual is responding *to* can provide the response with a standard of correctness.

For example, recall the case of Gloria and her poor exam score. When Gloria hides her score, so that no one else can see it, this seems to be a fitting response to the shamefulness of her score. But suppose now that an eccentric billionaire offers Gloria a large sum of money to respond to the shamefulness of her score by advertising it to the entire class. Gloria now has a reason to engage in this odd sort of action-response. But no sum of money could make such a response a fitting response to the *shamefulness* of her score. Although Gloria may have a reason to engage in this response, this is the wrong kind of reason. It is a reason for engaging in the response that does not contribute to the *fittingness* of the response. If Gloria is to engage in an action-response to the shamefulness of her score, only the score's shamefulness, intuitively, provides the right kind of reason for responding; in particular, the shamefulness of Gloria's score provides her with a reason to keep the score hidden from view, not to advertise it to the entire class.

Similarly, when Fahad runs away from the bear, this seems to be a fitting response to danger. But now suppose that an eccentric billionaire offers Fahad a large sum of money to respond to the danger posed by the bear by running recklessly towards it. It is true that Fahad now has a reason to engage in this very action-response. But once again, the billionaire's offer seems to be a reason of the wrong kind. For no sum of money could make running recklessly

towards the bear a fitting response to the *danger* posed by the bear. Only danger seems to provide the right kind of reason for responses to danger.¹⁷

My suggestion here is that action-responses can be supported by the wrong kind of reason in much the same way that mental states such as beliefs and desires can be. Indeed, the reasoning I have used in the preceding examples perfectly mirrors the familiar way in which beliefs and desires are often shown to be susceptible to the wrong kind of reason. To illustrate, suppose that an eccentric billionaire offers you a large sum of money to believe some falsehood or to desire a saucer of mud.¹⁸ Intuitively, you now have a reason to believe this falsehood or to desire the mud in question. But just as intuitively, the billionaire's offer is the wrong kind of reason to believe or to desire. It is a reason that does not contribute to the *fittingness* of believing a falsehood or of desiring mud.

So action-responses and mental states such as beliefs and desires are susceptible to the wrong kind of reason in much the same way. And it is for this reason, I suggest, that both mental states and action-responses are assessable for fittingness. Now, action-responses also possess another important earmark of susceptibility to the wrong kind of reason. Schroeder (2012) points out that it is typically difficult, perhaps impossible, to have an attitude for the wrong kind of reason, whereas it is easy to have an attitude for the right kind of reason. It is difficult, for example, to desire a saucer of mud because of a financial incentive, but it is easy to

¹⁷ On this note, some philosophers have tried to distinguish the right kind of reasons from the wrong kind of reasons by making a distinction between "object-given" reasons (reasons derived from the object) and "state-given" reasons (reasons to hold a particular attitude toward the object). See, for instance, Parfit (2001, 2011, Appendix A) and Piller (2006). However, this effort has met with resistance, in part because it is possible to recast properties of an attitude as properties of its object (e.g. the fact that I have an incentive to desire the saucer of mud means that the *mud* is such that, were I to desire it, something good would happen). For discussion, see Olson (2004), Rabinowicz and Rønnow-Rasmussen (2004), Lang (2008).

¹⁸ See Crisp (2000).

desire a saucer of pudding because of its desirability. Of course, if a billionaire offers you a large sum of money to desire a saucer of mud, there are steps you can take that may indirectly lead you to desire the mud. For example, you can attend more closely to the mud than you ever did before and search specifically for desirable characteristics. I doubt you will ever find any. But, the thought goes, if this process does lead you to form a desire for the mud, then you will have come to desire it for the right kind of reasons – the features of the mud. It is hard to imagine, and it may be impossible, to come to desire the mud *directly* upon hearing the billionaire's offer.

At any rate, action-responses also exhibit this earmark of wrong-kind-of-reason susceptibility. For example, consider Fahad and the bear. It is difficult for Fahad to engage in action-response to *danger* for any reason other than danger. After all, insofar as Fahad's behavior is a response to the billionaire's offer, it is not any longer an action-response to *danger*, specifically. Of course, if the billionaire offers Fahad a large sum of money to respond to danger by running recklessly towards it, there are steps that Fahad could take that might indirectly lead him to engage in the desired response. For example, Fahad could examine the danger and search specifically for characteristics of the danger that merit running recklessly towards it. I doubt that he could ever find any. But if he did, if he did come to respond directly and specifically to the danger posed by the bear by running recklessly towards it, this would be for the right kind of reasons – the features of this danger. Again, it is hard to imagine, and it

may be impossible, for Fahad to respond directly and specifically to the *danger*, with reckless running, immediately upon hearing the billionaire's offer.¹⁹

Now, the preceding analysis depends on the distinction between actions and action-responses. Gloria *does* have a reason to advertise her poor score, and this reason does not seem to be a reason of the wrong kind until we assess Gloria's behavior as a response to the shamefulness of her score. When we assess Gloria's behavior as an *action-response*, we assess it in relation to a circumscribed consideration (i.e. shamefulness) that generates fittingness conditions. Interestingly, if we assess Gloria's behavior as a response to the *billionaire's offer*, then we have a different standard of correctness. In particular, we find that if Gloria publicly announces her poor grade, this *would* be a fitting response to the attractive offer made by the billionaire. The fact that Gloria's score is shameful is a reason for her not to announce her grade, but no matter how shameful her score is, announcing her grade would still be a fitting response to the billionaire's offer. Therefore, if we assess Gloria's response in relation to the offer, then the shamefulness of her score becomes a reason of the wrong kind!

This suggests that reasons of the wrong kind naturally arise whenever we assess a response in relation to a circumscribed standard.²⁰ Since action-responses can be assessed in

¹⁹ Cf. Schroeder (2010, 2012), who claims that *some* actions can be done for the wrong kind of reason, but that they do not possess this typical earmark of wrong-kind-of-reason susceptibility, and Rabinowicz and Ronnow-Rasmussen (2012), who seem to claim that actions more generally can be done for the wrong kind of reason, but that they do not possess this typical earmark of wrong-kind-of-reason susceptibility.

²⁰ Schroeder (2010, p. 32) has made this point before: "So what does it take to give rise to a Wrong Kind of Reasons type set of issues? I think it takes very little. My view is that Wrong Kind of Reasons issues arise in any domain which is governed by a standard of correctness." One of my ambitions in this paper is to expand on Schroeder's claim by showing just how ubiquitous and available such standards of correctness can be. We find such standards whenever we engage in the ordinary assessment of responses.

relation to such circumscribed standards, action-responses can be supported by the wrong kind of reason. And, I suggest, this is all it takes for a response to be assessable for fittingness.

Now, it might be objected that there is a disanalogy between action-responses and mental states such as beliefs. Beliefs are often thought to be governed by an *internal* standard of correctness.²¹ By contrast, action-responses do not seem to be governed by internal standards of correctness. The examples I have given suggest that we evaluate action-responses in light of *external* standards, such as those generated by danger, shamefulness, or a billionaire's offer.²²

I grant this disanalogy, but I suggest that this has no effect on our ability to assess action-responses for fittingness. Fittingness assessments require only that a response can be assessed in light of a circumscribed standard, in a way that is insensitive to extraneous reasons. It does not matter whether the standard is internal or external to the response. As I have just argued, action-responses seem to be assessable for fittingness in a way that is familiar and intuitive. And action-responses can be supported by the wrong kind of reason. The fact that the circumscribed standards for action-responses may be external therefore does not seem to pose any obstacle for fittingness assessments.²³

²¹ For discussion of the view that beliefs are inherently governed by a norm of truth, or constitutively aim at the truth, see Williams (1973), Railton (1994), Wedgwood (2002), Boghossian (2003), Gibbard (2005), Shah and Velleman (2005), Steglich-Peterson (2006, 2009), Vahid (2006), and Whiting (2010, 2012).

²² It is *danger*, for instance, that generates a standard of correctness for action-responses to danger, not anything internal to responses to danger themselves. One can respond to a danger by running away from it, but one can also respond to a danger by embracing it. Only one of these responses is a fitting response to danger, and this has to do with the nature of danger, the object of the response, not with the response itself. By contrast, in the case of belief, it is intuitively not the object of the belief, but the belief itself that generates a standard of correctness. This is why beliefs are thought to be governed by an internal standard of correctness.

²³ Of course, philosophers are welcome to restrict their use of 'fittingness' to encompass only those fittingness phenomena that are associated with internal standards of correctness. But it is important to recognize that this

If all of this is right, then we are now in a position to achieve a deeper understanding of why the recognitional view of fittingness is false. I think it has been tempting for philosophers to think that the reason that mental states can be fitting or unfitting is that they represent the world in some way. But I suggest that this is not the fundamental explanation. The deeper reason that mental states can be assessed for fittingness is that they can be assessed in relation to a circumscribed standard, and so can be supported by the right kind of reasons. Indeed, since many mental states are thought to have an *internal* standard of correctness, these states are naturally assessable for fittingness. But other types of responses can still be assessed as fitting or unfitting, since they can be assessed in light of external, circumscribed standards. I have argued that the non-representing components of emotional responses can be assessed for fittingness in this way. And I have argued that action-responses can be assessed in this way as well. I conclude that the recognitional view is mistaken. The recognitional view inappropriately emphasizes the importance of correct mental representations for fittingness.

4. Elaboration and Difficult Cases

We now have two problem cases, then, suggesting that we should abandon the recognitional view of fittingness, as well as a deeper explanation for why it is false. First, against the recognitional view, the fittingness conditions for emotional responses seem to go beyond correct representation. Second, action-responses can be assessed for fittingness, yet this

would be a *revision* of our ordinary concept of fittingness. As I have been arguing, we have clear, pre-theoretic intuitions that responses can be fitting in relation to external standards. And such fittingness phenomena behave just as we would expect them to – in particular, they allow for a distinction between the right and the wrong kind of reasons. If we revise our concept of fittingness, we might miss this important similarity amongst fittingness phenomena. It would be much better, it seems to me, to understand fittingness as it really is: a quite general phenomenon that depends on circumscribed standards of correctness, whether internal or external.

cannot be explained by the recognitional view either. The deeper explanation for why the recognitional view is false is that fittingness conditions are generated by circumscribed standards quite generally, not specifically by standards for mental representations. We therefore need a more expansive conception of fittingness.

4.1 An alternative conception

A simple solution would be to join the recent push by some philosophers to accept fittingness as normatively primitive, even more primitive than reasons or values.²⁴ If fittingness is normatively primitive, then there is of course no obstacle to accepting the conclusion, for instance, that action-responses can be fitting, if this is what the evidence suggests.²⁵

But a more familiar reasons-first approach to the normative is also compatible with the expansive conception of fittingness that I advocate.²⁶ A fitting response can be understood simply as a response that is supported by the right kind of reasons. Of course, a reasons-first advocate will want to understand the “rightness” of the right kind of reasons without appealing to a notion of correctness. But so long as one does this in a way that is compatible with the fittingness of both emotions and action-responses, the resulting view will have an important advantage over the recognitional view of fittingness. According to this alternative view, both

²⁴ See Chappell (2012), McHugh & Way (2016), Howard (Forthcoming). See Rowland (2017) for a recent criticism.

²⁵ Indeed, these so-called “fittingness-first” advocates have already suggested the possibility that fittingness assessments may be extended to choices (see Chappell, 2012) and “perhaps, to [actions]” (see McHugh & Way, 2016, p. 586), though this conclusion has not been carefully defended.

²⁶ For examples of the reasons-first approach, see Scanlon (1998), Schroeder (2008), Skorupski (2010), and Parfit (2011).

action-responses and emotions are fitting when they are supported by the right kind of reasons.

It is beyond the scope of this paper to decide between fittingness-first and reasons-first approaches to the normative. In what follows, my goal will be to work out some of the difficulties and answer some of the questions that naturally arise for the position I have been defending.

4.2 Internal standards and mental representations

Consider first the case of belief. Because beliefs have an *internal* standard of correctness, we can assess them for fittingness without specific reference to an external standard of correctness. A belief that p is fitting *simpliciter* when it is the case that p . Indeed, the truth or falsehood of $\lceil p \rceil$ is the only consideration that bears on whether the belief that p is fitting simpliciter. And so, strictly speaking, only the fact that p is the right kind of reason for the belief that p to represent the world as it does – at least insofar as we are assessing the belief in relation to its internal standard.

But we can also assess beliefs as *responses*, in relation to *external* standards of correctness. Because beliefs aim at the truth, they are merited by certain external conditions. In particular, *evidence* provides the right kind of reason for believing. When a belief is supported by the evidence, by the right kind of reasons provided by the evidence, the belief is a fitting response *to the evidence* – or, as we typically say, warranted.

The two types of assessments mentioned here can of course come apart. A belief may be fitting *simpliciter*, but not a fitting response to the evidence; and a belief may be a fitting response to the evidence, but not fitting *simpliciter*.

We can extend this analysis to the mental representations of emotional responses, which also seek to correctly represent the world. These representations can be fitting *simpliciter* or fitting responses to evidence. But crucially, against the recognitional view, I have argued that the fittingness of the emotional response as a whole is not simply a matter of the fittingness of its mental representation. A fitting emotional response correctly represents, but also correctly mobilizes the individual in response to the external condition being represented. In order to be fitting, the emotional response must be supported by the right kind of reasons.²⁷

4.3 External standards and fittingness as a distinctively narrow assessment

I have argued in this paper that fittingness assessments are not limited to beliefs or emotional responses. We can also assess action-responses for fittingness. In assessing an action-response in relation to a circumscribed, external standard, we assess it in a way that is distinctively narrow.

Now, philosophers commonly distinguish between *pro tanto* justification and *all things considered*, or *all-in* justification. But I wish to suggest that fittingness assessments are even narrower than your typical, unqualified *pro tanto* assessment. An action-response may be *pro*

²⁷ Can the emotional response as a whole be warranted? Perhaps, but one would need to identify the value in relation to which such a response would be fitting. It is not clear to me that *evidence of danger*, for instance, provides the right kind of reason for the *whole* fear response, though it may. A better candidate might be *likely danger*. Perhaps fear is warranted when it is a fitting response to likely danger (even if there is no actual danger present).

tanto justified, yet fail to be a fitting response to some particular value or consideration. This is because an action-response can be *pro tanto* justified for the wrong kind of reason. Consider again the case of Gloria and her poor exam score. If the billionaire offers Gloria a large sum of money to advertise her poor grade to the class, then Gloria is *pro tanto* justified in doing so. But just as we concluded before, if Gloria advertises her poor grade, this would not be a fitting response to the score's shameful nature. Advertising the score may be *pro tanto* justified, but it is *pro tanto* justified by the wrong kind of reason.

If we wish to describe fittingness in terms of *pro tanto* justification, then, we might say that a fitting action-response is a response that is *pro tanto* justified by the right kind of reasons. If Gloria hides her shameful score, this is a fitting response to the score's shameful nature – it is *pro tanto* justified *by this shameful nature*.

4.4 The nature of external, circumscribed standards

I have said that responses can be assessed in relation to external, circumscribed standards that generate fittingness conditions. But how exactly do we determine the fittingness conditions generated by a particular standard?

Consider a simple example first. Why should we think that Fahad's running away from the bear is a fitting response to danger, or to that particular bear? Well, if we take a fittingness-first approach, we might simply claim that this is a first-order normative judgment, no more mysterious than the judgment that the hungry bear is a reason to run away. But even if we do prefer a fittingness-first approach, it is often helpful to confirm our intuitions by recasting the fittingness judgment in terms of reasons. And this, of course, will be the natural approach for

the reasons-first advocate. Consider: intuitively, the danger Fahad faces, the bear he is running from, provides him with the right kind of reason to run away. This confirms our intuition that his running away is a fitting response to danger, to the bear.

Consider next a more difficult case. What is a fitting response to the threat of nuclear war? Once again, we might approach this question by asking ourselves what reasons we take to be provided by this threat. Perhaps the threat of nuclear war provides us with the right kind of reason to anxiously construct an underground shelter. Doing so would, intuitively, be a fitting response to the threat of a nuclear attack. But suppose you believe that no shelter you could build would protect you, that there's indeed nothing you can do to save yourself. So, depressed, you go to the bar to drink your sorrows away. Is that a fitting response to the threat of nuclear war? This is tricky – is it the threat of nuclear war itself that provides you with a reason to despair, or is it the fact that there's nothing you can do about the threat of nuclear war that provides you with this reason? If the latter, then your despairing is a fitting response to the fact that there's nothing you can do about the threat of nuclear war.

In short, we should not be surprised that it is sometimes difficult to determine the fittingness conditions generated by a particular circumscribed standard; this difficulty parallels the difficulty of determining what reasons are provided by some particular consideration. Nevertheless, it is often helpful to test or confirm our intuitions by converting fittingness judgments into judgments about reasons, or vice-versa.

4.5 The nature of responses

The next difficulty has to do with the nature of responses. What is it that licenses us to regard some action or some physiological change as a *response*? In virtue of what is Gloria's action a response to the shamefulness of her score? In virtue of what is Fahad's elevated heart rate a response to the danger posed by the bear? Well, in both of these cases, we find that the action or the physiological change in question was causally downstream to a *representation*. Gloria's action was mediated, at least in part, by a representation of her score's shamefulness. Fahad's elevated heart rate, similarly, was mediated by some representation of the threat posed by the bear. I think this mediation by a representation is necessary for regarding an action or physiological change as a response to some consideration.²⁸ When such mediation is absent, the process in question cannot be regarded as a response. For instance, if Gloria folds her exam paper over absentmindedly, completely unaware that it contains her score, her action can hardly be construed as a response to her score's shamefulness.²⁹

Now, recall that on my view, any response can be assessed for fittingness, because any response can be assessed in relation to an external, circumscribed standard of correctness, given by what the response in question is a response *to*. Gloria's response to the shamefulness

²⁸ But of course, the correctness of the representation in question is not what determines the fittingness of the downstream response, as we have seen. The upstream representation simply makes it the case that the downstream change *is* a response.

²⁹ In the case of representations that are not veridical, I think we qualify response attributions in various ways. If Misa hallucinates a dragon and then runs away, we might say that her running was a response to a *hallucinated* dragon. Similarly, if Misa runs away from a tiny, harmless spider that she takes to be very dangerous, we might say that she was responding to what she *thought* was danger. Or, we might in such a case fall back to a representation that we can actually agree with, such as Misa's representation of there being a spider present. We might say that Misa's running away was an unfitting response to this tiny little spider. But it was certainly not a response to genuine danger.

of her score can be assessed for fittingness, because her response can be a fitting or unfitting response to the shamefulness of her score. Similarly, Fahad's elevated heart rate, a physiological response to danger, can be assessed for fittingness, because it is either a fitting or unfitting response to danger.

We are now in a position to say something about the practical scope of fittingness assessments. In order to endorse an action, for example, as a fitting response to some consideration, the action must actually be a response to that consideration. If Gloria advertises her poor score to the entire class, completely unaware of the billionaire's offer, then she is not responding to the billionaire's offer, and so it would sound odd to endorse her action as a fitting response to that offer.³⁰ But we are often interested, I think, in criticizing actions precisely by pointing out that they are not fitting responses to some possibly important consideration ("that was *not* a fitting response to how nice your sister has been to you lately!"). And this means that the reach of fittingness assessments can be very great. Your reading this paper, intuitively, is not a fitting response to the threat of nuclear war. The fact that it wasn't meant to be a response to that threat does not undermine the point of the criticism.

4.6 Difficult Cases

With this understanding of fittingness assessments in hand, we can now discuss some difficult cases where intuitions about fittingness are often unclear. As I have already noted, difficulty often arises when it is challenging to determine what reasons are provided by a particular consideration. In such cases, clarity can often be achieved by thinking carefully about

³⁰ Though perhaps it would be alright to say that Gloria's action happened to be fitting or correct in *relation* to the billionaire's offer.

the consideration in question, as well as nearby considerations that may also be generating fittingness conditions.

Consider the case of the counterproductive response. Suppose that Fahad, having escaped from the bear, is now up against a dragon with very sharp claws. But this dragon is quite intolerant of those who seek to escape his wrath. Any attempt at avoidance is met with an even more dangerous, unhappy dragon. Is running away a fitting response? Well, the dragon's claws are very dangerous, *but* running away would result in even greater danger. The conjunction of these two considerations, then, is a reason to stand still. So, we can say that keeping still would be a fitting response to the following conjunctive fact: that the dragon's claws are sharp, *but* that running would result in even greater danger.

However, you might have the intuition that running away would be a fitting response to the fact that the dragon has sharp claws – ignoring, as it were, the part about the dragon's wrath. We can justify this intuition by pointing out that, setting aside the dragon's wrath, the sharpness of the dragon's claws is a reason to run. In forming this intuition, perhaps we imagine a default or paradigm situation where it's just Fahad and the dragon's claws, without any unusual additional details.³¹

³¹ Of course, whether some consideration is a reason for a specific action always depends on *some* background features of the situation. For example, the fact that the dragon has sharp claws is a reason for Fahad to *run* only if he *can* run, if there is room to run, etc. It is important not to conflate these relevant background features with extraneous or wrong-kind reasons. Wrong-kind reasons or features are those that do *not* affect the reasons provided by a particular consideration. If an eccentric billionaire offers Fahad a large sum of money to run recklessly towards the dragon, for instance, then Fahad has a reason to do so. But this reason does not affect or determine the reasons provided by the dangerousness of the dragon (though it may outweigh these reasons).

My point in the main text is that when we assess a response in relation to a certain consideration in isolation, we seem to assume a certain (default or paradigm) background that helps determine what reasons are given by this consideration.

Counterproductive cases such as this are difficult, because the considerations being responded to seem to interact. As I have illustrated in the example above, I think the best way to approach such cases is either to acknowledge the interaction by assessing the action-response in relation to the conjunction of interacting considerations, or, where this is not desirable, to abstract away from the interactions and to assess the action-response in relation to one of the conjuncts in isolation. This approach seems reasonable to me, and it also seems to track very well the kinds of intuitions we have in these cases.³²

Difficulty can also arise in cases that involve competing, but overlapping values. Some activities, for example, are both dangerous and thrilling. When thrill-seekers run excitedly towards such dangers, is their response fitting or unfitting? Well, I think we should say that the thrill-seekers exhibit a fitting response to the *thrilling*. The fact that some activity is thrilling is a reason to pursue it. But the fact that some activity is dangerous is not by itself, intuitively, a reason to pursue it, for danger is something to avoid. Of course, certain activities are thrilling in part because they are dangerous. But the two values are nevertheless distinct and merit different responses. For example, when the thrill-seekers don their protective equipment, this is a fitting response to the dangerousness of the activity they are about to pursue. But their pursuit of this activity in the first place is a fitting response to the fact that it is thrilling – not,

³² It might be thought that a reasons holist (e.g. Dancy, 2004) has a different approach available. A reasons holist might argue, for instance, that in the dragon case, keeping still is a fitting response to the dragon's sharp claws (because in this case, the sharp claws provide Fahad with a reason to keep still). This position is compatible with my larger ambition in this paper, but I think even the holist should reject it. If Fahad, in keeping still, is really responding only to the dragon's sharp claws, and not to the part about the dragon's temper, then there's something quite odd about his response. Why is he keeping still? Intuitively, a fitting response to just this aspect of the situation would be to run. The holist can accept this conclusion by embracing the second of the two strategies I mentioned above.

intuitively, to the fact that it is dangerous. In summary, I think we should handle this type of case by teasing apart the competing values.³³

Similarly, it might be wondered whether it is fitting to be amused by an offensive joke. Well, if offensive jokes can be funny, then I think we should say the following: amusement is not a fitting response to the offensiveness of the joke, but it *is* a fitting response to the funniness of the joke. I think this is what we should say even if the joke is funny *because* it is offensive.

4.7 Meta-normative questions

One might worry, at this juncture, about disagreement in taste. If I find some joke funny that you find rather dull, is there a fact of the matter about which one of our responses is fitting, if any? I have no view on the matter. My goal in this paper is to convince you that fittingness assessments apply across a wide range of responses. Whether there are any objective fittingness facts, or whether there are relative fittingness facts, etc., is an interesting meta-normative question, just as it is an interesting question whether there are objective values or reasons. This issue is beyond the scope of this paper.

Nevertheless, one might wonder if on the expanded conception of fittingness that I advocate, *every* response will be fitting in relation to *some* circumscribed standard. And would that be a problem? Well, I think it is true that we can always invent a putative circumscribed

³³ But my larger points about fittingness are once again compatible with the position of the reasons holist. If the holist wishes to insist that the fact that something is dangerous is sometimes a reason to pursue it, then she can do so. But once again, even the holist may wish to agree that if it is the danger alone that is being responded to, and not any interacting considerations, the danger is something to avoid (see previous footnote).

standard in relation to which a response will be fitting, but I do not think that this is a problem, for such ad-hoc standards will not generally be of interest to us. For instance, suppose that Lydia has a bizarre attraction to sharp objects. Every time she sees a knife, she feels a strange desire to feel its sharp edges. In order to express herself to others, she might coin the term 'sharposity' to refer to the alluring feature that she takes sharp objects to possess. We might then say that when Lydia excitedly approaches sharp objects, this is a fitting response to their sharposity, though it may not be a fitting response to their dangerousness. But this would be to "buy in" to Lydia's point of view, and we may not wish to do that. We may simply wish to deny that sharposity is a legitimate or interesting standard for fittingness assessments.

Conclusion

The capacity to assess responses for fittingness is an important component of our normative psychology. But the nature of fittingness assessments has not been well understood. In this paper, I have argued that fittingness assessments have a distinctively narrow focus; fittingness assessments are possible whenever we can assess a response in relation to a circumscribed standard. This means that fittingness does not have anything specifically to do with correct mental representations, and so we should reject the recognitional view of fittingness. In particular, I have argued that the fittingness conditions of emotional responses go beyond correct representation. And, I have argued that action-responses can be assessed for fittingness, a phenomenon which cannot be explained by the recognitional view. As a result, I have suggested that we should embrace a more expansive conception of fittingness, one that embraces a wider range of responses.

Chapter II: Emotion Regulation and Evaluative Understanding

Introduction

Anxiety can consume us, rouse us, and motivate us to worry about some future threat. Such worrying may lead us to consider the extent to which the threat really does imperil us, and as a result, we may *discover* some aspect of the threat that we did not recognize before. Similarly, sadness can overtake us, slow us down, and motivate us to ruminate about some loss or setback. Such rumination may lead us to consider the extent and implications of our misfortune, and thereby, we may learn something about our misfortune, our loss or setback, that we did not previously appreciate. Anger, likewise, can engulf us, inflame us, and motivate us to simmer about some injustice we've faced. Such simmering may lead us to consider more closely the nature and severity of the injustice, and as a consequence, we may learn something about the injustice that enriches our understanding. And so on for many other emotions – shame, happiness, envy, guilt, fear. Emotions, through their ability to direct our attention and our thinking, can guide us to a better understanding of the true nature and value of their objects; emotions, by alerting us to the presence of significant facts, can promote our

evaluative understanding.¹ This is a point that has been made in one way or another by philosophers across a range of philosophical disciplines.² I shall discuss these views in Section 1.

Of course, it is a commonplace that emotions can also blind us, bias us, and bend our wills in ways that are not conducive to evaluative understanding. It is therefore tempting to think that even if emotions can sometimes guide or direct us towards greater evaluative understanding, *ideal* evaluative understanding will only be achieved if we guide or direct the emotions themselves – that is, if we properly *regulate* our emotions.³ But this hypothesis raises an interesting and underexamined philosophical question: *how* should we regulate our emotions in order to achieve evaluative understanding? What forms of emotion regulation should we embrace, and what forms of emotion regulation should we avoid? It is the aim of this paper to answer these questions and to describe a practical deliberative model that we as individuals can rely upon to decide which emotion regulation strategies we should use and *when* we should use them, if our goal is evaluative understanding.

In section 1, I shall describe the range of philosophical views endorsing the idea that emotions can enhance our understanding by directing our attention and thinking. This will set the stage for the subsequent discussion of the importance of emotion regulation for evaluative

¹ I borrow the term ‘evaluative understanding’ from Michael Brady (2013), who provides one of the clearest articulations of the relationship between emotions and evaluative understanding. I shall focus on evaluative understanding in this paper, as opposed to non-evaluative understanding, because I believe that the relationship between emotions, emotion regulation, and evaluative understanding is particularly rich and easy to misunderstand. But much of what I will have to say applies to non-evaluative understanding as well.

² In particular, I shall discuss views from the philosophy of emotion, epistemology, and feminist philosophy. See Section 1.

³ Cf. Brady (2013). Brady argues that emotions enhance our evaluative understanding when our attention is virtuously controlled. But I shall be emphasizing that certain forms of emotion regulation can *themselves* enhance evaluative understanding – it is not just the emotion virtuously controlled that enhances understanding.

understanding. In section 2, I will describe some of the different forms of emotion regulation. I will begin with an influential classification scheme developed by the psychologist James Gross, but I will also introduce a new distinction, a distinction between emotion regulation strategies that lead us to *engage* with objects of evaluation and emotion regulation strategies that lead us to *disengage* from objects of evaluation. In section 3, I will consider a tempting theory, the *Engagement View*, according to which engaged forms of emotion regulation can contribute to evaluative understanding, whereas disengaging forms of emotion regulation hinder the development of evaluative understanding. However, I will reject the Engagement View as untenable. I shall argue that *disengaging* forms of emotion regulation are actually vital for the development of evaluative understanding. But this point will raise a further, difficult question: when should individuals rely on engaging forms of emotion regulation, and when should they rely on disengaging forms of emotion regulation? In Section 4, I shall attempt to answer this question by proposing a practical deliberative model that will help individuals to decide which form of regulation is appropriate in a given context, if the goal is evaluative understanding. In Section 5, I will defend the model I propose against objections.

1. Emotions and Evaluative Understanding

Let us begin by examining how emotions themselves can enhance evaluative understanding. First, it is important to note that evaluative understanding requires more than the simple recognition of certain concerns or values – certain aspects of evaluative reality. Evaluative understanding requires, in addition, that we grasp the *nature* and *significance* of

these concerns or values.⁴ Anxiety, for instance, may correctly represent some future possibility as a threat, and so may amount to a recognition of this threat. But in order for such anxiety to enhance our evaluative understanding, the anxiety must help us to grasp *why* the future possibility is a threat, to what *extent* it is a threat, how difficult it will be to cope with the threat, and so on. In general, then, although emotions may provide us with initial evaluations that allow us to recognize important concerns (e.g. threats, losses, injustices), this is not yet enough for evaluative understanding. Emotions must also help us to grasp the nature and significance of the concerns which they help us to recognize. And indeed emotions characteristically do enhance our evaluative understanding in this way. How do emotions do this?

Well, an important observation, one that is made by a number of philosophers writing about the emotions, is that emotions, beyond providing us with an initial evaluation, draw our *attention* to or make *salient* certain facts that are relevant to the initial emotional evaluation.⁵ Allan Gibbard, for example, writes that “Emotions involve patterns of salience: if I am angry at you I dwell on special features of what you have done; if what you did instead had made me afraid then I would have dwelt on different features.”⁶ Catherine Elgin, similarly, writes that “Emotions direct attention. They reveal certain aspects of a domain as worthy of notice.”⁷ And

⁴ My position here is neutral on whether understanding is to be explained in terms of knowledge. For early discussion on the relationship between understanding and knowledge, see Zagzebski (2001), Kvanvig (2003), and Grimm (2006). For a review of the literature on understanding, see Baumberger, Beisbart, and Brun (2017).

⁵ See, e.g, Rorty (1980), De Sousa (1987), Roberts (1988, 2003), Narayan (1988), Jaggar (1989), Gibbard (1990), Hookway (2003), Greenspan (2004), Lance and Tanesini (2004), Elgin (2008), Goldie (2008), Brady (2013).

⁶ Gibbard (1990), p. 136.

⁷ Elgin (2008), p. 44.

De Sousa seems to think that this power of emotions to make certain facts salient is *indispensable* for the development of understanding. According to De Sousa, “No logic determines *salience*: what to notice, what to attend to, what to inquire about ... The role of emotion is to supply the insufficiency of reason ... an emotion limits the range of information that the organism will take into account.”⁸ If De Sousa is right about this, then emotions are a necessary part of any complete account of understanding.

But as Michael Brady points out, emotions do not simply direct our attention and then let go – emotions also tend to keep our attention *fixed* on concerns related to the emotional object. That is, emotions do not only direct us to engage with the emotional object and related concerns, they also typically make it difficult to “disengage our attention and shift it elsewhere.”⁹ As Brady puts it, emotions “*capture* and *consume* attention.”¹⁰ As a result of this attentional syndrome, emotions tend to propel inquiry in a certain direction. As Hookway remarks, “emotions can have a role in regulating the paths taken by our reflection and deliberation ...”¹¹ Brady offers a helpful example, asking us to consider the feeling of being awakened by loud sounds, downstairs, in the middle of the night.¹² In such a situation, we are likely to experience fear or anxiety, and our attention will prod us in a certain direction: we might perk up our ears, searching for confirmatory evidence of danger downstairs, and our mind might begin to race through the possible explanations of the downstairs commotion.

⁸ De Sousa (1987), p. 191-195.

⁹ Brady (2013), p. 92.

¹⁰ *Ibid.* All of this is well in line with empirical conceptions of emotion and its effect on attention. See Yiend (2010) for a review of the literature.

¹¹ Hookway (2003), p. 80. See also Morton (2010).

¹² Brady (2013), p. 95-96, p. 102.

The relationship between emotions and understanding now begins to emerge. Beyond providing us with an initial evaluation, emotions can also mobilize directed reflection. By guiding our attention and sending us down certain paths of investigation, emotions can propel us to discover important, relevant facts that we might not have appreciated otherwise. As I suggested in my opening paragraph, anxiety can help us to learn something new about a future threat, sadness can help us to learn something important about a loss or setback, and anger can help us to better understand the nature of an injustice that we've suffered. In general, by propelling us to discover facts that are relevant to our emotional concerns, emotions can enhance our evaluative understanding: they can lead us to update our initial emotional evaluations in light of a better grasp of the concern in question.^{13,14} Worrying about a future threat, for instance, may lead one to recognize that the threat is actually worse than one had imagined, or perhaps not as bad as initially feared. The individual's understanding of the *importance* of the threat, the danger posed by it, is thereby enriched. Similarly, sadness can motivate us to reflect on the implications of a loss or setback, thereby propelling us to better understand the significance of the misfortune in question. Anger, likewise, can deepen our understanding of the pervasiveness and severity of a certain form of injustice.

This insight has been appreciated by feminist philosophers. Uma Narayan, for instance, argues that the emotional reactions of "insiders," or victims of oppression, can lead them to

¹³ This type of view is most clearly articulated in Brady (2013).

¹⁴ Evidence from experimental psychology suggests that we are constantly evaluating the stimuli that we encounter and the things that we think about (Zajonc 1980, 1985). Even in seemingly calm moments, we form such affectively charged evaluations. And so, whenever we attend to some non-evaluative fact, we immediately evaluate its significance. In particular, if the non-evaluative fact is about some evaluative object, we immediately evaluate the significance of this fact for the evaluative object in question.

develop a richer understanding of oppression than even the most sympathetic amongst outsiders:

The victim [insider] ... may feel a complex and jumbled array of emotions: anger at the perpetrator, a deep sense of humiliation, a sense of being 'soiled' by the incident, momentary hatred for the whole group of which the perpetrator is a part, rage at the sort of history that has produced and sustains such attitudes, anger and shame at one's powerlessness to retaliate, a strong sense of solidarity with those who face the same problems, and maybe even pity for the stupidity of the perpetrator.¹⁵

Emotions guide the insider's attention and thinking in certain directions, enriching the insider's understanding of the nature of oppression, its pervasiveness, and its severity. Emotions deepen, that is, the insider's evaluative understanding. Alison Jaggar makes a similar point:

Only when we reflect on our initially puzzling irritability, revulsion, anger or fear may we bring to consciousness our 'gut-level' awareness that we are in a situation of coercion, cruelty, injustice or danger. Thus, conventionally inexplicable emotions, particularly though not exclusively those experienced by women, may lead us to make subversive observations that challenge dominant conceptions of the status quo.¹⁶

In this way, emotions can promote our moral understanding. But the mechanism by which this happens is of course not unique to the moral sphere. Emotions can also guide us to richer prudential and aesthetic understanding. Disappointment with some aspect of an activity – a hobby, or a career, for instance – can prompt reflection that may lead an individual to a deeper understanding of the activity's true value. Perhaps collecting stamps gets old after a while, one realizes, an instance of diminishing returns; perhaps grading papers is less exciting than one had hoped, involving more tedium and repetition than anticipated. Similarly, one's unexpected joy and engagement with a science fiction novel may prompt reflection about the

¹⁵ Narayan (1988), p. 39.

¹⁶ Jaggar (1989), p. 167.

nature of science fiction, and eventually, a deeper appreciation for the genre's merit. In summary, the relationship between emotions and evaluative understanding seems to hold across a wide range of values.

Of course, emotions do not *always* contribute to evaluative understanding. Indeed, the very same attentional syndrome that enables emotions to promote evaluative understanding also seems capable of leading us astray. By focusing our attention on some limited domain, some particular aspect of the emotional object (e.g. how threatening it is), emotions can lead us to miss or ignore important considerations that, in a cooler moment, we would agree are relevant to our evaluative understanding. As Peter Goldie puts it, emotions “[skew] the epistemic landscape”.¹⁷ If our goal, therefore, is evaluative understanding, we shall need to make room for emotion regulation.

But it may be that not all forms of emotion regulation contribute to evaluative understanding. In the next section, I shall introduce some of the different forms of emotion regulation.

I. Emotion Regulation Strategies

The most influential classification scheme for emotion regulation strategies is the one developed by psychologist James Gross.¹⁸ Although we need not accept every detail of Gross's view, his model of emotion regulation will serve as a useful (and well-regarded) starting point.

¹⁷ Goldie (2008), p. 159.

¹⁸ Gross (2014).

In the previous section, we discussed how emotions, once generated, can guide our attention and our thinking. But we have not yet said anything explicit about how emotions are generated in the first place. Yet this is important, because Gross's model of emotion regulation is built directly on top of a more basic model of emotion *generation*. The idea is that once we understand the different components of emotion generation, we can classify the different emotion regulation strategies based on the components of emotion generation that they target. Let us begin, then, with Gross's model of emotion generation.

Gross calls his model of emotion generation the "modal model," because it posits features and processes that psychologists widely agree are important to emotion generation. The model includes four components. 1) The first is the *situation* in which the individual finds herself. The thought is that emotions do not arise from a vacuum, but are responses to the changing features of the individual's situation.¹⁹ For example, if Sasha is being evaluated by her supervisor, she is likely to experience different emotions than when she is lounging on her sofa, watching Netflix. However, 2) in the same situation, different individuals may pay attention to different aspects, and this may result in divergent effects on emotional responding. And so, *attention* is the second component of the emotion generation process. If Sasha attends closely to her supervisor's criticisms, but pays little attention to the supervisor's praise, her emotional response will be affected accordingly. Nevertheless, 3) individuals who are in the same situation, and who attend to the same stimuli, may still experience different emotions. This is because different individuals may interpret the meaning or significance of situational features

¹⁹ Gross is careful to note that 'situation' is to be understood very broadly, so as to include imagined situations capable of eliciting an emotional response (Ibid, p. 5).

in different ways. Therefore, the third component of emotion generation is *appraisal*, a process which refers to the interpretation of the situational features to which one attends.²⁰ If Sasha interprets the supervisor's praise as empty or disingenuous, she will not feel good about it, however much she attends to it. Finally, 4) the result of all of this is the emotional *response* itself, which includes changes in feeling and motivation, physiological changes (such as changes in emotional expression, muscle tenseness, and respiration), and also the cognitive, attentional changes that were of interest to us in Section 1. Sasha may now dwell on what she perceives as a negative review and may be motivated to ruminate about the quality of her work. And of course, this process of directed attention will lead to further appraisals, reigniting the emotion generation process.

For each component of the emotion generation process, Gross identifies a class of emotion regulation strategies targeting that component. For instance, 1) situations can be *selected* or *modified* with the aim of shaping emotional experience. Fahad, for example, always has a lousy time at parties, so he avoids them (*situation selection*). Or if he does go to a party, he modifies the situation by stealthily turning down the obnoxiously loud music (*situation modification*). 2) Attention can be regulated by engaging in *attentional redeployment*, which can involve distracting oneself, averting one's gaze, or otherwise refocusing one's attention away from the source of emotional concern. Mindfulness and meditation practices are widely used forms of attentional redeployment. 3) Appraisals can be regulated through *cognitive change*, such as reappraising a situation or thinking positive thoughts. Fahad, initially

²⁰ We need not take appraisal to be a conscious process. We also need not take appraisal to be a process that always requires the possession of concepts.

disappointed with the grade he received on his exam, might reappraise the situation by reminding himself that he is taking the class pass/fail, and that the grade doesn't really matter anyway. Cognitive-behavioral therapies make wide use of cognitive change strategies. 4) Finally, *response modulation* usually refers to the regulation of the non-cognitive aspects of an emotional response (e.g. relaxing one's muscles); the cognitive aspects of the emotional response can be understood as belonging to a new emotion generation process, and their regulation is covered by attentional redeployment and cognitive change strategies.

It's worth noting that many emotion regulation techniques make use of more than one of Gross's emotion regulation categories (meditation techniques, for example, often make use of both attentional redeployment and physiological response modulation). Nevertheless, Gross's model provides us with a reasonable initial framework for exploring the relationship between emotion regulation and evaluative understanding. With Gross's model in hand, we can now introduce what I take to be a philosophically crucial distinction: the distinction between emotion regulation strategies that allow us to better *engage* with evaluative objects and emotion regulation strategies that lead us to *disengage* from evaluative objects.

As I have noted in Section 1, emotions characteristically mobilize reflection about the nature and significance of an emotional concern. Anxiety, for instance, leads us to think about the nature and significance of a future threat, sadness leads us to think about the nature and significance of a loss or setback, and so on. Let us say that the individual who is in the grip of such directed reflection is *engaged* with the emotional concern. Now, certain forms of emotion regulation regulate emotional experience while keeping the individual engaged with the emotional concern. By contrast, other forms of emotion regulation operate by leading the

individual to *disengage* from the emotional concern, precluding further reflection on its nature and significance. This distinction is important, because as we have seen, engagement with an emotional concern allows the individual to uncover considerations bearing on its nature and significance, and this can promote evaluative understanding. But we have also noted that emotional reflection, since it is guided in a certain direction, characteristically uncovers only a narrow subset of considerations bearing on the emotional concern. Thankfully, however, there are certain forms of emotion regulation, such as reappraisal, that can keep the individual engaged with the emotional concern, while helping the individual to uncover new, possibly *countervailing* considerations that are relevant to the emotional situation. As a result, such forms of emotion regulation can play an important role in enhancing evaluative understanding.

Let us consider this point about reappraisal in more detail. As I have already noted, emotional reflection characteristically draws our attention to a narrow subset of considerations bearing on the value of the emotional object. But when we reappraise a situation, we look at it in a different way – we reinterpret the situation in light of different considerations. Suppose, for instance, that you were planning to meet a friend for the night, but at the last minute they have to cancel. Naturally, you're quite disappointed, and you begin ruminating about how fun the night would have been – what a pity. But then you reappraise the situation. You recognize that at least now you'll have some extra time to catch up on that novel you've been reading. The night is not lost after all! And you'll have another opportunity to meet your friend on a different day. So the whole thing is not a big deal.

In this way, by reappraising the situation, you uncover different, possibly countervailing considerations bearing on the value of your emotional concern. Of course, we often reappraise

situations in this way simply to make ourselves feel better.²¹ But importantly, the considerations that we draw upon when we reappraise are typically considerations that we would independently endorse. The reappraisal wouldn't make us feel any better if we did not take it to capture a genuine feature of the evaluative situation.²² And so, by reappraising a situation, we broaden our view. We move beyond our initial emotional reflection, often challenging or modifying our initial evaluations. By remaining engaged with the emotional concern, but regulating our emotional reflection through reappraisal, we are able to more fully consider the nature and significance of the evaluative situation. But of course, just as emotional reflection can lead us astray by leading us to focus on a narrow subset of relevant considerations, reappraisals can be problematic if they lead us to minimize or ignore the legitimate considerations initially brought up by emotional reflection. And so in order to ensure that our reappraisals genuinely enhance our evaluative understanding, we should ideally *integrate* the considerations brought up by our reappraisals with the considerations brought up by our emotional reflection. We should bring all of these different considerations together and

²¹ More generally, reappraisal efforts are activated by emotion regulation goals - by the desire, for instance, never to be excessively angry, or by the desire to avoid distress.

When an emotion regulation goal is present, reappraisal efforts can occur automatically and effortlessly; a sudden awareness of one's anger may lead one to reappraise the offender's intentions as possibly benign. But reappraisal efforts can also be initiated in a way that is deliberate and effortful. Grappling with shame and embarrassment, an individual may challenge herself to consider whether she should care so much about what other people think, leading to thoughtful reflection and reappraisal. Both types of reappraisal - effortless and deliberate - have been identified in the literature on emotion regulation. See, for instance, Gyurak, Gross, and Etkin (2011). The more effortful form of reappraisal can itself come in different types. For instance, some deliberate reappraisals are prepared ahead of time; the individual resolves to reappraise a future stressful situation in a certain way. But reappraisals often need to be generated on the fly; for instance, certain forms of cognitive-behavioral therapy involve learning how to challenge cognitive distortions as they arise, thereby reappraising the emotional situation (e.g. Beck's cognitive therapy, 1967, 1976, or Ellis's Rational Emotive Behavior Therapy, 1997).

²² Though as we shall see a bit later, we can often fool ourselves by selectively attending to positive considerations and ignoring negative ones.

give each its proper weight in order to build a more comprehensive picture of the evaluative situation.²³

Consider an example. Suppose that Maria is incredibly anxious about the speech she has to give tomorrow. Her mind keeps running through the various nightmare scenarios: what if she forgets all of her lines? That would be so awkward. What if she is simply unable to answer one of the Q&A questions? That would be terrifying. Maria feels lousy about all of this, and so she decides to reappraise the situation. Is it really *likely* that she will forget all of her lines? Maria reminds herself that she is very well-prepared, having rehearsed her presentation dozens of times by now. So yes, it is possible that she will forget her lines, but this isn't likely. And her friends will be in the audience – they'll be a source of comfort. Worst comes to worst, she can always pull out her paper and just start reading from it. And if she does have an awkward moment – which, admittedly, she might – so what? Would that really be the end of the world?²⁴

²³ A more precise characterization of such integration would require telling a story about the various ways in which considerations interact to yield a conclusion about what to believe. Yet this appears to be no easy task, as some philosophers worry that reasons can do more than just compete with one another (see, for instance, Dancy, 2004, Nair, 2016, and Cullity, 2018). It is beyond the scope of this paper to address this issue. For my purposes, an individual integrates two considerations X and Y when she takes into account both considerations and how they interact before settling on a conclusion. For example, the individual might compare the weights of X and Y. Or, she might determine that X somehow affects the significance of Y as a consideration. Or, the individual might discover that, in light of consideration X, Y does not seem to be as relevant a consideration as initially suspected.

²⁴ It might be objected that nothing is the end of the world but *the end of the world*. Can it really enhance our evaluative understanding to recognize that our personal concerns are not literally the end of the world? The answer is *yes*: the point of this idiom is not to remind us that our concerns do not matter. Rather, the point of the idiom is to put our concerns in *perspective*. Maria gains evaluative understanding by reminding herself that if she does forget all of her lines, this would be bad, but not nearly as bad as her anxiety makes it out to be.

But there is a potentially deeper worry here. There are certain evaluative perspectives according to which much of what we typically care about isn't actually worthy of our concern. The Stoic may claim that much of what we are anxious about is beyond our control and therefore irrational to fear. A Buddhist may claim that many of the things that concern us are impermanent, that our attachment to these things is a source of suffering. The account that I shall develop doesn't rule out any of these views *a priori*. Rather, my account simply assumes that, as agents, we do not begin from a position of complete evaluative understanding. We do not, that is, begin with any fixed

In reappraising her situation, Maria remains engaged with the object of her anxiety. And though her reappraisals broaden her view and uncover new considerations, she is careful to *integrate* these new considerations with the considerations originally brought up by her anxiety. She acknowledges that it is possible that she will forger her lines, but now recognizes that this is not likely, given how much she's practiced. And she accepts that she might have an awkward moment, but she ultimately questions whether this would be as bad as she initially felt it would be. As a result of this process of reappraising the situation and integrating these reappraisals with her initial concerns, Maria seems to enhance her evaluative understanding of the situation.

Of course, as I have already noted, reappraisals do not always enhance evaluative understanding. Consider an alternative scenario in which Maria is not at all prepared for her speech. Once again, she anxiously envisions the nightmare scenarios. But then, in an effort to make herself feel better, she begins to reappraise. It occurs to her that she's given speeches before, and they've all gone quite well. Latching onto this thought, Maria says to herself, "Hey, I'm sure it's going to be *fine*! I shouldn't worry so much." Maria reaches this new evaluative conclusion by reappraising the situation, to be sure, but her reappraisal involves very limited

view, whether Stoic or Buddhist. Instead, we begin with an incomplete evaluative understanding, and with enough epistemic humility to be willing to *reflect* on the value of things. I will argue that engaging and disengaging forms of emotion regulation can both play a role in enhancing our evaluative understanding. For all I know, the use of these emotion regulation strategies, coupled with reflection, will eventually lead us to something like the Stoic or Buddhist view. But we might just as easily be led in a quite different direction. So it makes sense for Maria to begin with her personal concerns, to reflect on the extent to which these concerns really matter. It could be that she is systematically mistaken, and that one day she will become a Stoic or Buddhist and abandon many of these concerns. But Maria does not yet know where true wisdom lies, and so the best she can do is rely on strategies that seem to be enhancing her evaluative understanding.

My thanks to Dan Jacobson for helpful discussion on this point.

engagement with the object of her emotional concern. In particular, Maria fails to integrate the new consideration with the concerns brought up by her anxiety. Yes, she has given speeches before, but she's also been far more prepared in the past than she is this time around. The fact that she has given speeches before should not by itself, then, quell her concerns about failure. But since Maria is in this case primarily motivated by a desire to feel less anxious, she does not take the time to integrate her reappraisal with her initial anxious concern. Instead, Maria latches onto the positive reappraisal, quickly concludes that she will be fine, and moves on. In this way, Maria uses her reappraisal to facilitate a premature *disengagement* from her emotional concern. As a result, it does not seem that Maria's reappraisal promotes her evaluative understanding.

Importantly, however, this is not what happened in the original example. In the original example, although Maria was motivated to regulate her anxiety, she was also motivated to reappraise the situation in an engaged fashion, to determine the true nature of the future threat. And this did seem to promote her evaluative understanding. My claim so far, then, is that reappraisals tend to enhance evaluative understanding when they are part of a larger process of thoroughly *engaging* with an emotional concern, which typically requires integrating the considerations brought up by reappraisals with the considerations brought up by one's initial emotional reflection.

Consider now a different worry. It might be objected that Maria could not have gained evaluative understanding by reappraising her situation, because in order to accurately reappraise, she must have already possessed the relevant understanding. But this is not true. Let us suppose that prior to reappraising, during her anxious reflection, Maria either believed

that she faced a serious threat or had no clear view on the matter. In an effort to reappraise, Maria asked herself questions with the aim of more closely evaluating the seriousness of the threat: how well prepared am I? How bad would it really be if I forgot a few lines? Upon reappraising, upon reflecting on these new questions, Maria arrived at conclusions that changed her evaluative understanding in the right direction. Cheered by her evaluation of her own preparedness and of the tolerable badness of forgetting her lines, Maria reappraised the future threat as not-so-serious. Now, of course, Maria may have had to rely on her previous understanding when she evaluated her own preparedness and the badness of forgetting her lines. But this is a familiar feature of reflection. We work our way to new understanding by building upon previous understanding. In this case, Maria arrives at a piece of new understanding: the threat she faces is not so serious. Maria gains new evaluative understanding by reflecting on new questions and by making new inferences.²⁵

²⁵ What about a case where Maria already *does* understand that the future threat is not serious, but feels anxious anyway? There are two possibilities. If the emotion is amenable to reappraisal, Maria might remind herself of what she already understands in order to bring her emotional appraisals in line with her evaluative understanding. In such a case, reappraising would not serve to enlarge the understanding, but would serve to bring emotional experience into agreement with the understanding. The second possibility, though, is that the anxiety is recalcitrant. Despite her best efforts to reappraise the situation in light of her genuine evaluative understanding, Maria might continue to feel anxious. In such a case, reappraisal would not only fail to enlarge the understanding, it would fail to bring emotional experience in conformity with the understanding. This explanation relies on the now-familiar point that emotions, unlike judgments, are characteristically susceptible to recalcitrance – emotions can resist contradiction. For further discussion, see D’Arms and Jacobson (2003).

But if one’s emotions do not *initially* reflect one’s evaluative understanding, does one *really* understand? Doesn’t the incorrect emotional reaction reveal one’s lack of understanding? Not necessarily. Consider the common aversive reaction to insects. Most of us understand quite well that the typical moth or spider cannot hurt us and cannot infect us with anything. Nevertheless, we are wired to react aversively to small, gray-black creepy-crawlies. This doesn’t demonstrate that we lack understanding; it only shows how emotions, by shifting our cognition in characteristic ways, can momentarily prevent us from accessing understanding that we really do possess. Consider: we do not say of an exhausted mathematician who is failing to solve straightforward proofs that she has lost her mathematical knowledge, or that she doesn’t understand how to complete these proofs. We say that the exhausted mathematician understands how to complete these proofs, but that she is too tired to do so right now – too tired to access her own understanding. Similarly, emotions, by shifting our cognition in characteristic ways, can temporarily prevent us from accessing and acting on our understanding. But a recognition that our emotion is unfitting may come soon after, along with reappraisals grounded in our genuine evaluative

Reappraisals, then, by promoting thorough engagement with an emotional concern, can enhance evaluative understanding. But as we have seen, reappraisals can also facilitate inappropriate and premature disengagement from an emotional concern, precluding reflection on its nature and significance.

But if such disengagement is problematic, then we might also think it questionable to rely on emotion regulation strategies that are disengaging by nature.²⁶ Many forms of meditation, for example, invite the practitioner to redirect attention away from the current evaluative object and towards some other object, such as the breath. This form of regulation might help individuals to feel calmer and happier, but one might worry about whether it hinders the development of evaluative understanding. Now, some forms of mindfulness do allow the practitioner to be aware of emotional thoughts, but the practitioner is encouraged *not* to engage with these thoughts by judging, evaluating, or considering. Instead, the practitioner is encouraged to non-judgmentally *note* the passing emotional thought, and then, perhaps, to let the passing thought go. Disengaging forms of emotion regulation, then, seem to preclude the kind of focused reflection that can promote evaluative understanding.

understanding: “it’s just a moth – it can’t hurt me.” All of this is not to deny, of course, that *sometimes*, perhaps often, emotional reactions do reveal a lack of evaluative understanding. It’s just that this relationship is not necessary. We should also accept that our emotional reactions can sometimes cast doubt on our previous, genuine understanding, so that by reappraising we can actually reconfirm our understanding.

²⁶ Many forms of emotion regulation are not inherently engaging or disengaging, and so their contribution to evaluative understanding depends on the context. Drinking alcohol, for instance, may facilitate either engagement or disengagement. In this paper, I shall mainly be focusing on those forms of emotion regulation that are clearly engaging or disengaging, in order to illuminate the relationship between this distinction and the development of evaluative understanding.

2. The Engagement View

With this distinction between engaging and disengaging forms of emotion regulation in view, an intriguing thesis may suggest itself. In particular, we might suspect that whereas engaging forms of emotion regulation can enhance evaluative understanding, disengaging forms of emotion regulation necessarily detract from evaluative understanding, or at least hinder its development. Call this the *Engagement View*.

The Engagement View seems fairly plausible on a first pass. After all, it's hard to see how one could gain understanding about some evaluative object without engaging it, without reflecting on its nature and significance. Similarly, it is difficult to see how disengaging from some evaluative object could possibly lead to a gain in evaluative understanding.

The Engagement view is further bolstered by an important clarification. The Engagement view does not claim that trying to adopt a *distanced* or *detached* perspective on some evaluative object is bad for evaluative understanding. On the contrary, so long as the individual remains engaged with the evaluative object, a more distanced or detached perspective may make it easier to reappraise the evaluative situation. So the Engagement View is compatible with the advice of Adam Smith, who encourages us to evaluate our deeds and the deeds of others by adopting the perspective of a fair and impartial spectator.²⁷ Similarly, the Engagement View is compatible with Hume's idea that in order to overcome biases in our evaluations, we should adopt a "common point of view,"²⁸ a point of view that discounts our

²⁷ Smith (1759), e.g. III.I.2.

²⁸ Hume (1739), Book 3, Part 3, Section 1. Hume points to biases that arise from seemingly irrelevant factors, such as our distance from the evaluated object or its familiarity.

personal biases. These perspectives are compatible with the Engagement View insofar as they call for detached or distanced engagement rather than disengagement.²⁹ In order to properly challenge the Engagement View, one must show that genuine disengagement can promote evaluative understanding.

I can think of three ways in which disengaging emotion regulation strategies can enhance evaluative understanding. The first relies on the idea that every activity has an opportunity cost. By engaging with a certain evaluative object, the individual loses the opportunity to engage with a different, possibly more important object. The other object may not only be more important, it may also be less well understood than the evaluative object that is currently being engaged. And so, the individual who fails to disengage from the original object and re-engage with the second object loses out on important evaluative understanding. For example, suppose that Dalya is returning from her first day of work at her new job. The day didn't go as well as she had expected. The job is going to be way more demanding than she had anticipated, Dalya realizes, and it looks like she'll have to learn a new programming language on her own time. Dalya ruminates about this during her commute back home, and she begins to form plans that will help her adjust to the new situation. She's still in the grip of all of this when she arrives at home and finds that her partner is completely distraught. Recognizing that a new,

²⁹ For a contemporary view, see Kross, Ayduk, and Mischel (2005), Grossman and Kross (2014), and Kross and Ayduk (2008, 2011, 2017). According to Kross and colleagues, *self-distancing* allows individuals to process emotional concerns thoughtfully, without becoming enmeshed in intense negative affect that could lead to narrow emotional reflection and rumination.

In my view, reappraisals can be achieved in a number of different ways, and self-distancing is one of these. By self-distancing, we can identify novel, relevant considerations that we failed to identify from a more "self-immersed" perspective. But there are other ways to generate reappraisals. For example, the individual can search for considerations that challenge her initial evaluation. If the individual begins with a negative evaluation, the result is likely to be a positive reappraisal. It currently does not appear that these different implementations of reappraisal have different regulatory efficacy (see Ranney, Bruehlman-Senecal, and Ayduk, 2017).

more important, and less well understood evaluative object has just presented itself, Dalya takes a moment to regulate her own emotions. She takes a deep breath and reappraises her own troubles as not important right now. That is, she implements disengaging forms of emotion regulation. This allows her to engage the new evaluative object, her partner's distress, with a clear mind. In other words, disengaging forms of emotion regulation are crucial in situations where evaluative understanding is best promoted by redirecting attention away from one evaluative object and towards another.

The second way in which disengaging forms of emotion regulation can enhance evaluative understanding relies on the idea that we can learn about evaluative objects by making comparisons. In particular, disengaging from one evaluative object and re-engaging with another can help an individual to gain insight about the *original* evaluative object. This is because engaging with a second evaluative object can trigger a novel reappraisal of the original evaluative object, a reappraisal informed by a *comparison* between the two objects of evaluation. Consider, for example, the stressed-out employee, Lucas, who is incredibly nervous to find out whether he has earned that important promotion. There is only so much that Lucas can learn by engaging with this evaluative object, by thinking about it and its properties. But then Lucas goes to the park. He sits by the river. He sees it flowing by, hears the sounds of the water washing past him. How soothing – how beautiful! Attending to this spectacle of nature, Lucas finds a kind of peace, and he says to himself: “Does any of this really matter anyway? Isn't *this* what matters? Here I am, able to appreciate life and the world. Earning that promotion isn't going to change any of it. Do I want that promotion? Sure I do. But this is not as big a thing as I was making it out to be. It's like a drop of water compared to this river.” By disengaging

from thoughts about his promotion and engaging instead with the beauty of the river, Lucas is able to discover a fact about his original concern (the promotion) that he might not otherwise have easily appreciated. In particular, by engaging with the river's beauty, Lucas begins to appreciate the *relative* unimportance of his promotion. This insight is made possible by a comparison with a second evaluative object. So disengaging from one evaluative object and engaging with another can enhance evaluative understanding by enabling, upon re-engagement with the original object, a comparison-based reappraisal.

The third way in which disengaging forms of emotion regulation can enhance evaluative understanding also involves disengagement followed by re-engagement. But this time, the reason that disengagement helps is not because it facilitates a comparison, but simply because it reduces emotional biases.³⁰ As I explained in Section 1, emotions make certain facts salient and direct our thinking in certain directions. This means that emotion regulation strategies that involve reflection (e.g. engaged reappraisals) may themselves be subject to these biases. In particular, the attentional biases present in the context of an intense emotion may hinder even sincere attempts at reappraisal, since the emotion will keep pushing attention in a certain direction. Additionally, emotions seem capable of making their evaluative objects appear more important than they really are, perhaps by funneling attention away from considerations that would put the emotional object in perspective. Unfortunately, these issues may be difficult for engaged forms of regulation to address. The individual may make a respectable effort to

³⁰ Mindfulness meditation, for instance, has been shown to reduce a number of affect-related biases, such as the negativity bias, the correspondence bias, and the sunk cost bias, among others. See, for instance, Kiken and Shook (2011), De Raedt et al (2012), Garland and Howard (2013), Hafenbak et al (2014), Lueke and Gibson (2015, 2016), Golubickis et al (2016), Tincher et al (2016), and Hopthrow et al (2016).

examine the different sides of the evaluative object, but if the individual is under the influence of an intense emotion, certain sides of the object may seem excessively salient, and the importance of the object may be exaggerated. Now, these impressions may be difficult to overcome or ignore, and so even well-meaning attempts at engagement may lead to biased results. But if the individual begins by relying on a disengaging form of emotion regulation, such as mindfulness meditation, the individual may be able to clear her mind in a way that facilitates more accurate reappraisals upon re-engagement.³¹

For example, recall the case of Maria, who is anxious about her speech tomorrow and is worrying about the nightmare scenarios. She tries to reappraise, but she ends up returning to her anxious concerns. So she stops herself, takes a deep, full breath, and attends to the feeling of air rushing through her nostrils, inflating her chest. She slows herself down and continues to focus on her breath. After some time, perhaps she redirects her attention to the rest of her body, then, slowly, to her position in space, then to the sounds around her, and so on. As she becomes more mindful of her environment, the salience of her previous concern diminishes. Now it is her breath, her position in space, and the sounds around her that are salient. As a result, Maria's anxiety diminishes as well. When Maria re-engages with her anxious concerns, she is calmer, and her attentional biases have been redirected. Whereas before her mental space was consumed by her anxious concerns, now she is able to adopt a more distanced, detached perspective. It is easier for her to look at the situation objectively and to consider new

³¹ There is evidence that mindfulness facilitates cognitive reappraisals. See, for instance, Garland et al (2011), Troy et al (2013), Hanley and Garland (2014), Hanley et al (2014), Garland et al (2015a), and Garland et al (2017). This has led to the development of the Mindfulness-to-Meaning theory, which holds that the salutary effects of mindfulness on well-being are due to its facilitation of positive reappraisals. For discussion, see Garland et al (2009) and Garland et al (2015b, 2015c).

perspectives. And perhaps the concern doesn't seem as all-encompassing as it did before. All of this makes it easier, we can imagine, to reappraise the concern accurately. The idea here should be a familiar one. Sometimes it's enough to simply take a deep breath before thinking through some stressful issue. This minor disengagement is sometimes enough to clear the mind, to downregulate the emotion and its biases. And all of this can facilitate reappraisal upon re-engagement. In this way, disengagement can help to promote evaluative understanding.

These, then, are three ways in which disengaging forms of emotion regulation can help enhance evaluative understanding. Disengaging forms of emotion regulation can help us to manage opportunity costs, to make comparisons between evaluative objects, and to reduce emotional biases. In light of these considerations, the Engagement View seems untenable, since it claims that disengaging forms of emotion regulation detract from or hinder the development of evaluative understanding. Of course, disengaging forms of emotion regulation *can* be detrimental to evaluative understanding, when they lead us to simply ignore or minimize important concerns. But the upshot of this section is that disengaging forms of emotion regulation can also play a central role in enhancing evaluative understanding. This raises an interesting normative question: when should individuals rely upon engaging forms of emotion regulation, and when should they rely upon disengaging forms of emotion regulation? This is the topic of the next section.

3. When to Engage, When to Disengage

In section 2, I identified three ways in which disengaging forms of emotion regulation can enhance evaluative understanding: disengaging forms of emotion regulation can help us to

minimize opportunity costs, to make important comparisons between evaluative objects, and to reduce emotional biases. In this section, I will draw on these benefits of disengagement in order to develop a practical deliberative model for emotion regulation that can tell us when we should engage and when we should disengage, if our goal is evaluative understanding.

A close examination of the benefits of disengagement allows us to divide the reasons for disengaging into two general categories: 1) disengaging from an evaluative object can help to enhance our understanding of some *other* evaluative object, henceforth some *independent* evaluative object. We saw this in the case of Dalya, who had to disengage from her personal concerns in order to learn about the nature of her partner's distress, which represented an independent concern. 2) disengaging from an evaluative object can also help us to enhance our evaluative understanding of *that very object*, henceforth the *original* object. Disengaging from the original object can paradoxically enhance our understanding of it by facilitating important comparisons between this original object and other, important evaluative objects (as in the Lucas case). Disengaging from the original object can also enhance evaluative understanding by helping the individual to reduce emotional biases, allowing for better re-engagement with the original object (as in the Maria case).

In summary, then, disengaging forms of emotion regulation can enhance our evaluative understanding of independent evaluative objects as well as original evaluative objects. Therefore, when we are in the grip of an emotion, wondering whether to engage or disengage, we should consider these two general reasons for disengagement, if our goal is evaluative understanding.

If we want to be methodical, we should begin by making sure that the object of our emotional concern, the original object, is really the evaluative object we should be considering. Dalya, for instance, though she was initially concerned with her own problems, noticed an independent evaluative object of great importance – her partner’s distress. Recognizing that this independent evaluative object was more important to understand and less well understood than her original concern, Dalya decided to disengage, with the aim of better understanding this independent evaluative object.

Let’s examine the claim I just made more closely. I said that the distress of Dalya’s partner was both *more important* to understand and also *less well understood*. These two factors together contribute to the opportunity cost for Dalya of ignoring her partner’s distress. But if her goal is evaluative understanding, why shouldn’t Dalya simply focus on the evaluative objects that are least well understood? Why should she also prioritize the objects that are likely to be most important to understand? Well, so far, I have been speaking rather loosely of the goal of enhancing evaluative understanding. But to be more precise, our goal is to gain evaluative understanding of the things that *matter most* (at least initially). That is, our goal is not to maximize evaluative understanding in a purely quantitative way. If that were so, we might praise Dalya for ignoring her partner’s distress and diving into a tome of ancient Greek plays, since she knows even less about these than the likely cause of her partner’s distress! No, our goal should be to enhance our evaluative understanding of those objects that, based on our current understanding, are most likely to be important. That said, quantity matters too. If some independent concern is more important to understand, but already quite well understood, it might make sense to remain engaged with an original concern that is less important to

understand, but which is less well understood. Therefore, in deciding whether to disengage from an original concern and engage instead with an independent concern, we should consider both how important each object is *and* how well understood it is.

Of course, it is impractical to constantly deliberate about whether one should reflect on original or independent concerns. What we really want is for individuals to have a *sensitivity* to the considerations I have just mentioned. This sensitivity can perhaps be achieved in the way that habits and sensitivities are usually achieved – namely, by *consciously* practicing the activity that will eventually become an unconscious habit or sensitivity. For example, suppose that Dalya’s partner has criticized her in the past for focusing too much on her own concerns when a more serious, independent concern is present. Recognizing the truth in her partner’s words, we can suppose that Dalya formed a plan to more frequently consider whether she should shift her attention to certain independent concerns. Following through with this plan might eventually inculcate in Dalya a more effortless sensitivity to the importance of these independent concerns.

Let us now turn to the second reason for disengaging, which is to enhance our evaluative understanding of the original evaluative object. Suppose that we have already decided that the original evaluative object is the one that merits our attention. Still, there is a question of whether our evaluative understanding of this object would be best served by engaging or disengaging. How should we proceed? Well, at any given moment, we need some reasonable way of quickly comparing the prospect of disengagement to the prospect of engagement, or what we might call the *engagement-disengagement prospect*. We already

know that disengagement can help us to make important comparisons and to reduce emotional biases. But we also need to think about the benefit of engagement.

Engagement with an evaluative object is beneficial for evaluative understanding when it is *productive* – when it helps us to uncover new considerations and to integrate these into a broader, more comprehensive picture of the evaluative situation. Engagement with an evaluative object is *not* beneficial when it is *unproductive* – when it does not lead to any progress in constructing a more complete view of the evaluative situation. Engagement is often unproductive when it is *perseverative*, when engaged reflection about the object continuously rehashes and recycles old considerations. Perseveration is a well-known feature of certain emotional disorders, such as depression and anxiety.³² In these disorders, depressive rumination and anxious worrying sometimes keep the individual locked in an unproductive loop that continuously turns up the same negative considerations. But perseveration is a common product of emotional thinking in general, and is not limited to emotional disorders. When we are angry at someone, upset about some loss, or anxious about some future threat, we often find ourselves going over the same considerations again and again in our minds. This is neatly explained by the attentional biases of emotions discussed in Section 1. Since our emotions shift our attention in a certain direction, our thinking is likely to flow in a certain direction, and eventually this can result in rehashing the same considerations over and over. This is not to say that emotional biases play no role in enhancing our evaluative understanding. Initially, these biases can help us to discover new, relevant considerations bearing on the emotional concern

³² For an overview, see Nolen-Hoeksema, Wisco, & Lyubomirsky (2008).

(see Section 1). But emotions frequently stick around longer than necessary, as if insistent on making sure that their point has been made. And this can and does lead to unproductive perseverance.³³

So when one is initially confronted with an emotional concern, it seems quite reasonable to initially follow the emotion, to engage with the emotional concern, perhaps without regulation. This will often allow the individual to uncover some important considerations about the emotional object. Think about Maria and her anxious reflections about the nightmare scenarios. Although in Maria's case, further reflection revealed that these nightmare scenarios were unlikely, one can surely see the benefit of uncovering these worries in a situation where the nightmare scenarios are quite likely. But once Maria uncovered the relevant concerns, we can imagine that she began to perseverate, to continuously recycle these concerns over and over in her mind. At this point, emotion regulation is required in order to enhance evaluative understanding. Engaging forms of regulation may initially be appropriate. Cognitive reappraisal may help Maria to uncover new sides of the emotional concern while still remaining engaged. Insofar as this is productive, engagement is worthwhile. But suppose that Maria finds it difficult to reappraise effectively, since she is so affected by her anxious biases. Here disengagement can play an extremely important role. The appropriate form of disengagement depends in part on the intensity of Maria's anxiety. In some cases, a single deep

³³ Of course it may be that mulling over some consideration repeatedly can help one to more fully grasp this consideration and its relevance. I do not at all object to such mulling, but only to the unproductive extension of such mulling, which is perseveration. Can we introspectively distinguish productive mulling from unproductive perseveration? This is a difficult question, but I suggest that the two processes may *feel* different, at least in their extreme versions. We can feel certain considerations "sinking in" as we briefly mull them over. Similarly, if we are on the lookout for unproductive perseveration, we should be able to spot it ("I've already gone over this").

breath can suffice to reduce emotional biases. In other cases, a more extended period of mindfulness meditation (for example) may be necessary to ready one's mind for effective re-engagement.

At some point in this process, the individual will have developed a deep enough understanding of the emotional concern to begin putting it in perspective, to begin comparing it with other concerns. At this point, disengaging from the emotional concern and engaging with important, independent evaluative objects can serve to both regulate the emotion, reducing emotional biases, *and* to facilitate comparisons between evaluative objects. In my view, this is a critical step on the way to full evaluative understanding. Since emotions characteristically magnify the importance of their evaluative objects, emotional concerns must be put in perspective. The individual may remind herself of her loving family, or of the good things in life, the beauty of nature, etc. – the specific focus of comparison will of course vary from person to person and from situation to situation. But such comparisons play a critical role. We need to ground ourselves, evaluatively speaking, by keeping the importance of novel evaluative concerns in perspective. The relevant process of comparison can be rather extended, much like Lucas's trip to the river, but it can also be very brief, consisting in a rapid refocusing of one's attention. Maria, for example, might pause her anxious reflection in order to briefly think about her family, which she adores. In doing so, Maria might make it easier for herself to put her initial emotional concern into perspective, through a comparison-based reappraisal: Maria can remind herself that however the speech goes tomorrow, this is just a very small part of a big life filled with value.

The process I have been describing thus far is dynamic and fluid, but it is also principled. The individual begins by engaging with the evaluative object. Initially, this might just mean following the emotion's attentional promptings. But once it seems that emotional reflection is no longer productive, perhaps due to perseveration, engaging forms of emotion regulation, such as reappraisal, may allow for more constructive engagement with the evaluative object. But as soon as such regulated engagement seems unproductive, disengaging forms of emotion regulation become important, at first simply to reduce emotional biases, and eventually to make comparisons between the original concern and important, independent evaluative objects. In other words, the model I am advocating calls on us to "test-drive" engagement, to determine whether we should keep engaging, or whether we should disengage, by keeping track of how engagement is going (is this helping me to better understand the situation? Or, haven't I already thought about this?).

We can complicate this model by noting that perseveration is not the only way in which reflection can be unproductive. Sometimes emotional reflection is derailed by mind-wandering. Some consideration reminds us of something unrelated, and our thinking takes a detour. Whether we should re-engage with the original concern or continue engaging with the detour depends once again on how important and how well understood each object is. Engagement is also unproductive when it leads us to false conclusions, perhaps due to bias or improper motivation (see next section). But unless these biases or motivations are transparent to us, it's hard to see what we can do about them on our own. For this reason, it is often helpful to share our reflections with others, as this provides a useful check on our own biases and motivations.

Additionally, it is sometimes possible to predict the productivity of engagement even without test-driving it. Perhaps we already know that the evaluative object is quite well understood, for example, so that engagement will be unlikely to yield anything new. Or perhaps we feel that our emotion is so intense that engagement couldn't possibly be as productive as engagement preceded by at least some brief disengagement. In such cases, it of course makes sense to rely upon disengaging forms of emotion regulation even without test-driving engagement. But ultimately this is because in such cases we are able to calculate ahead of time the engagement-disengagement prospect.

Finally, we should say something about how this process comes to an end. When should we stop reflecting about our emotional concerns? Well, if our only goal were to enhance evaluative understanding, then reflection about the original emotional concern should only end when the opportunity cost for evaluative understanding becomes too high. But in reality, our goal of enhancing evaluative understanding is one goal among many, and our desire for evaluative understanding is often derived from our need to act. So typically, the reflection process described here should end when the individual has gained enough evaluative understanding to diminish the importance of gaining further understanding below the importance of acting. For example, once Maria determines that she is well-prepared for her speech and that the threat she was worried about is not very great, she may finally feel ready to put this understanding into action. Maria might, for instance, decide to rely upon a disengaging form of emotion regulation – not to further enhance her evaluative understanding, but to *relax*, to disengage from her present worries, and to wind down for the rest of the evening. On the other hand, if Maria had determined through her anxious reflection that the

threat she faced was severe, she may have decided, ultimately, that the best course of action is to stay up late preparing for the speech. In other words, in the typical situation, reflection should come to an end when the goal of achieving sufficient evaluative understanding to act has been realized.

Putting it all together, here is the practical deliberative model of emotion regulation that I propose, if our goal is evaluative understanding. First, we should disengage from our original emotional concerns when we judge that some independent evaluative object is more worthy of our attention, because of how important and how poorly understood it is. Once we've settled on an evaluative object to examine, we should typically test-drive engagement and look for productivity or perseveration in our reflection (though sometimes a simple prediction of productivity is possible; and sometimes we may encounter causes of unproductivity other than perseveration). When we encounter perseveration, we can initially disengage with a view to reducing emotional biases. Once we have a sufficiently clear picture of the emotional concern, we can disengage with a view to comparing the original object with important, independent evaluative objects, allowing us to place the importance of the evaluative object in perspective. Ultimately, we can stop reflecting once we have realized our goal of enhancing evaluative understanding.

This deliberative process, I suggest, is the best way to secure evaluative understanding that is both rich and deep.

4. Will It Work?

Earlier in the paper, I argued that both engaging and disengaging forms of emotion regulation can help promote evaluative understanding. But in the previous section I have advanced a stronger claim. I have claimed that if our goal is evaluative understanding, we *should* use engaging and disengaging forms of regulation in a particular, organized way. But of course, the claim that we should use this method to enhance our evaluative understanding invites more skepticism than the claim that this method *can* enhance evaluative understanding. In this section, I will strive to fend off such skepticism.

It might be objected that although the method I propose can enhance evaluative understanding, there are a number of reasons why the method might fail, and so it is not clear that we should advise people to use it. Looking at the different sides of an evaluative object and putting it in perspective may sound reasonable, on the one hand. On the other hand: individuals may lack the capacity to properly regulate their emotions or keep track of how engagement is going; their background evaluative beliefs may be faulty, resulting in unsound inferences; they may be improperly motivated, more interested in feeling good than in enhancing their evaluative understanding; they may stop deliberating prematurely, before all considerations have been uncovered and integrated; and finally, they may be inappropriately *swayed* by new considerations, and thereby led astray. In short, we have reason to be skeptical about whether the typical agent is really well-advised to adopt this deliberative method.

I shall address this objection in two steps: first with an analogy, and then more methodically. Here is the analogy. Suppose that Detective Lydia is trying to determine who committed some crime – was it suspect *A*, *B*, or *C*? Lydia might begin with a hunch that it was *A*,

because he looks a bit suspicious. But Lydia should not have *A* arrested on this basis alone, we think; she needs to develop a fuller understanding of what happened the night of the crime. She needs to gather evidence, to collect considerations for and against the hypothesis that *A* did it, as opposed to *B* or *C*. As Lydia collects evidence, she will likely begin to form a better picture of what occurred on the night in question.

Now, there is no *guarantee* that this process will improve her epistemic position. For one, Lydia may be a lousy detective; she might lack the capacity to find evidence effectively. But the solution to this problem is not to advise Lydia to refrain from collecting evidence, but to teach her how to do so more effectively, or to find another detective who can help her. It is also possible that Lydia's background beliefs about the case are faulty, so that collecting evidence will not improve her understanding of what happened. Of course, if we know this to be the case, we should bring this to Lydia's attention. But if we have no reason to suspect that Lydia's background beliefs are faulty, then it seems odd to advise her not to collect evidence, out of fear that her background beliefs *might be* faulty. Now, it could also be the case that Lydia is improperly motivated. Maybe she just wants the case to be over as soon as possible, so she can go home. This could indeed lead Lydia to investigate improperly. But again, the solution to this problem is not for Lydia to refrain from investigating. Lydia must regulate her improper motivations and collect evidence carefully, if she really wants to understand what happened the night of the crime. But isn't it possible that Lydia will be inappropriately swayed by some new piece of evidence? It is indeed possible. But once again, the solution to this problem is not to advise Lydia to refrain from investigating. In order to avoid being swayed, Lydia must strive

to keep in mind *all* of the considerations she has so far uncovered – she must use her regulation skills to avoid being swept up by a small set of new considerations.

Analogously, the deliberative method that I proposed in the previous section can fail to enhance an individual's evaluative understanding for a number of reasons. But in each case, the solution is not to refrain from using the method, since the kind of reflection it calls for seems instrumentally *necessary* for enhancing evaluative understanding of one's present emotional concerns; in order to understand the nature of some emotional concern, the individual must grasp the considerations for and against her initial emotional evaluation. To refrain from reflecting on these considerations is to give up on the goal of achieving full evaluative understanding. In what follows, then, I shall consider the primary ways in which this deliberative method can fail, and I shall suggest that in each case, the solution for the typical agent is not to refrain from using the method.

Lack of capacity. The first challenge arises for individuals who lack the capacity to implement the method properly. This may be due to lack of training in the required emotion regulation strategies, for example, or it may be due to a lack of cognitive ability. In the former case, if the individual's goal is evaluative understanding, the solution is of course to help the individual improve her regulation skills, not to advise her to refrain from deliberating. In the latter case, if the goal is to enhance the individual's evaluative understanding, then the solution is to assist the individual through the deliberation process in a way that is sensitive to the individual's ability. For instance, a parent can help a child to gain evaluative understanding of some emotional concern by assisting in the reappraisal process and generating relevant considerations. Insofar as we are *all* imperfect deliberators, we may all have a reason to seek

the help of others in this way, at least for difficult evaluative questions. This is fully compatible with endorsing the kind of deliberative method that I propose, since we can engage and disengage from evaluative concerns and reflect on them *together*.

Inaccurate background beliefs. Reappraisals depend upon background beliefs, some of which are evaluative (e.g. how bad it would be to forget one's lines) and some of which are non-evaluative (e.g. the likelihood of forgetting one's lines given that one has practiced). If these background beliefs are faulty, then reappraisals that rely on them are unlikely to enhance evaluative understanding. But this is no strike against the deliberative method that I propose. If we have reason to believe that some individual's background beliefs are faulty, we can bring this to his or her attention. But if there is no special reason to suspect faulty background beliefs, then it is odd to advise the individual to refrain from deliberating, on the off chance that background beliefs *might* be faulty (we wouldn't advise the typical agent to refrain from using *modus ponens* to make inferences from her background beliefs, for example, unless we had a special reason to think her background beliefs were faulty).

Improper motivation. One worry is that the typical individual is more interested in feeling good than in fully understanding her emotional concerns. That may well be true. But again, the solution is not to advise such an individual to refrain from deliberating at all. If the individual does have the goal of evaluative understanding (perhaps in addition to the goal of feeling good), then the individual must be mindful of improper motivation and regulate her thinking accordingly.³⁴ For example, the individual might ask herself whether she has reached

³⁴ Cf. Brady (2013).

some conclusion just because it would make her feel better. In general, for those individuals concerned about wayward motivations, this step could be part of the reappraisal process itself.

Incomplete consideration. The typical individual may fail to consider or integrate all relevant considerations related to some emotional concern, and so may arrive at faulty evaluative conclusions. This could be due to premature closure, in which the individual stops reflecting before all relevant considerations have been examined or integrated with one another.³⁵ Or the individual might be inappropriately *swayed* by a small subset of considerations. Whatever the precise details might be, the solution is not to advise the typical agent to refrain from deliberating. The solution is to warn the individual against such pitfalls and to propose specific methods for avoiding them. Premature closure can be avoided in part by making sure that one has completed every step of the deliberative method that I proposed (“have I thought through my emotional concern? Did I then reappraise, searching for opposing considerations? Did I try to put the emotional concern in perspective by engaging with an important, independent evaluative object?”). And to avoid being swayed, one of course needs to use the very emotion regulation strategies that I have recommended – engaged forms of regulation such as reappraisal can help the individual to keep the whole range of considerations in mind, and disengaging forms of regulation can help reduce emotional biases that underlie inappropriate *swaying*. An additional step that might help avoid swaying is to *summarize* from time to time the competing considerations that one has unraveled (“on the one hand, ...”). This

³⁵ Premature closure is a known problem in clinical medicine. Graber et al (2005) surveyed cases of diagnostic error and found that the most common cognitive contribution to error was premature closure.

can act as a check on any inclination to focus predominantly on any narrow subset of considerations.

In summary, the deliberative method that I propose does not guarantee evaluative understanding, and there are a number of reasons why such a process might go astray. Nevertheless, the solution to these problems is not to advise individuals to refrain from using the method, but rather to address the problems in question. Ultimately, if our goal is evaluative understanding, then we *should* use a method similar to the one that I propose.

Of course, it is important to recognize that proficiency with the kind of method that I have proposed may require a great deal of practice, perhaps even a lifetime of practice. In this way, the deliberative method is not like a set of rules which, if carefully followed, guarantees the right result. It is rather more like a set of skills, which, if carefully practiced, allow one to more deeply and thoroughly investigate the nature of emotional concerns.

Conclusion

In this paper, I have suggested that emotion regulation strategies can be divided into engaging and disengaging forms. Although it may initially appear that only engaging forms of emotion regulation serve to enhance evaluative understanding, I have argued that this is not the case. Disengaging forms of emotion regulation actually play a central role in the development of evaluative understanding. Although mindfulness meditation, for example, may initially seem to be bad for evaluative understanding, since it prompts us to disengage from and essentially ignore the evaluative object that initially concerned us, I have tried to show that

such disengagement, if implemented wisely, can actually play a central role in enhancing evaluative understanding.

In this paper, I have also proposed a deliberative model for deciding when to rely upon engaging forms of emotion regulation and when to rely upon disengaging forms of emotion regulation, if our goal is evaluative understanding. Though this model is not guaranteed to enhance evaluative understanding for every individual and on every occasion, I have suggested that this deliberative model, or one like it, is instrumentally necessary for evaluative understanding. Therefore, we should seek to overcome any obstacles to our implementation of this model. The alternative is to give up on our quest for emotional wisdom.

Chapter III: On the Distinctive Final Value of Negative Emotions

Introduction

A common idea among philosophers and laypersons alike is the thought that negative emotional experiences can be valuable for their own sake – that negative emotions can possess not only instrumental value, but also *final* value.¹ Even when negative emotions seem counterproductive, even when they are painful, these emotions can still be valuable, the thought goes, and it can still be the case that we ought to experience them. There seems to be something inherently valuable, for instance, about being *angry* at injustice, about *grieving* over a loss, and about feeling *guilty* when one has done something wrong. There is value, the thought goes, in being appropriately pained by the negative features we find in the world and in ourselves. And it is often claimed that a *calmer* response to these negative features would

¹ Although instrumental value is often contrasted with intrinsic value, its true opposite is final value. An object can be valuable for the sake of some other end (instrumental value), or it can be valuable as an end in itself (final value). By contrast, to say that an object is intrinsically valuable is to say that it is valuable in virtue of its internal (non-relational) properties. But some philosophers have argued that an object can possess final value without possessing intrinsic value (see, for instance, Korsgaard, 1983, Kagan, 1992, O'Neill 1992, and Rabinowicz and Rønnow-Rasmussen, 2000). I do not think that any of this matters for the arguments in this paper, but for the sake of precision, and to avoid any unnecessary commitments, I shall speak of the final value of negative emotions when characterizing my own position. I shall use 'intrinsic value' only when characterizing the views of others, if they use this expression, or if the phrase is particularly idiomatic in a given context.

Also, as shall become clear in what follows, it is not my view that all negative emotions are painful. Following Lazarus (1991), I take negative emotions to be those that register a "harmful person-environment relationship" (p. 22), or in philosophical terms, a negative value, such as injustice, personal loss, or personal wrongdoing.

somehow be lacking, somehow missing the point, somehow deficient in moral virtue. If this is right, then painful negative emotions, when appropriate, possess a final value that is distinctive, a value that calmer states cannot possess.

Amia Srinivasan puts the point well. Conceding for the sake of argument that anger is often counterproductive, Srinivasan argues that anger, when apt, is *intrinsically* valuable, and that this value is distinctive. It cannot be possessed by a calmer response:

I want to suggest that getting angry is a means of affectively registering or *appreciating* the injustice of the world, and that our capacity to get aptly angry is best compared with our capacity for aesthetic appreciation. Just as appreciating the beautiful or the sublime has a value distinct from the value of knowing that something is beautiful or sublime, there might well be a value to appreciating the injustice of the world through one's apt anger—a value that is distinct from that of simply *knowing* that the world is unjust.²

Though Srinivasan here contrasts *knowing* that the world is unjust with *anger* at the world's injustice, it is clear that she would mark a similar contrast between *calmly judging* that the world is unjust and being *angered* by injustice. To calmly judge that the world is unjust is not yet to "*feel* the ugly facts that structure our political reality."³ Anger, then, has a distinctive final value that calmer occurrent mental states cannot possess.

Troy Jollimore, writing on grief, argues that some forms of suffering are intrinsically valuable, since they provide us with a correct understanding of the world:

[A] failure to grieve, where grieving is appropriate, indicates a misunderstanding of the nature of the world one lives in ... a failure to comprehend the magnitude and significance of the loss that has occurred ... the fact that a person who ought to

² Srinivasan (2018, p. 132). Italics in the original.

³ Ibid, p. 141. Italics mine. Srinivasan claims that when anger is counterproductive, yet apt, individuals face a "substantive normative conflict" (p. 127). If calmly judging injustice to be present had the same final value as being angered by injustice, then presumably such conflicts could be easily avoided. Srinivasan would surely deny this.

experience grief fails to do so *itself* demonstrates that she is not correctly picturing her situation.⁴

For Jollimore, this is no contingent matter. Grief, on Jollimore's view, is no mere belief, but a distinctive cognitive state that necessarily involves suffering,⁵ and the elimination of which would prevent us from understanding a certain aspect of the world.⁶ Like Srinivasan, then, Jollimore would endorse the claim that there is a distinctive final value in appreciating loss through grief that is not to be found in the simple belief or judgment that a loss has occurred.

Hilary Bok, in a similar spirit, claims that the pain of guilt is essential to the recognition that one has done something wrong, comparing the pain of guilt to the pain of heartbreak:

[S]upposing that what one has lost was in fact love, and that one has indeed lost it, one can fail to suffer only if one walls oneself off from these facts ... Pain is not only an appropriate response to the collapse of love but the only response that accurately registers what has happened.⁷

Once again, this is no contingent matter. Adequate recognition of the fact that we have done something wrong requires us to be pained: "recognizing that we have acted wrongly necessarily and appropriately brings suffering in its wake."⁸ To calmly judge, then, that we have acted wrongly would not yet be to fully recognize our wrongdoing. Guilt and heartbreak, then, possess a distinctive final value that calmer mental states cannot possess.

Antti Kauppinen argues that anger is required for *both* the proper appreciation of important values and the expression of our concern for these values:

⁴ Jollimore (2004, p. 340-341). Italics in the original.

⁵ Ibid, p. 339.

⁶ Ibid, p. 342.

⁷ Bok (1998, p. 169).

⁸ Ibid.

The right kind of anger has value that is independent of its possible epistemic and motivational benefits ... If we eliminate anger, we fail to properly appreciate and may even lose personal relationships ... and moral relationships that depend on respect and mutual accountability.⁹

[S]uppose that I catch you in the act of trying to deceive me for financial gain. One possible response is to treat you as a nuisance to be managed ... But taking such an objective attitude, in Strawson's (1962) terms, is not consistent with respecting you. Respect requires instead that I'm willing to hold you accountable, circumstances permitting. If we continue to interact, I must show my anger and resentment ..."¹⁰

For Kauppinen, too, calmer responses will not do. "To respond properly to the value of people ... it is not enough if we reason and act in the right way."¹¹ Kauppinen, it seems, holds that anger possesses final value for both epistemic and moral reasons – anger is required for appreciating certain values and also for expressing our concern for these values. A calmer response would fail to properly register such values and express our concerns.

In short, a number of philosophers seem to endorse the following claim about the final value of negative emotions:

Distinctive Final Value (DFV): Painful negative emotions, when fitting, possess a distinctive final value, for epistemic or moral reasons, that calmer mental states cannot possess.

In this paper, I shall argue that DFV is false. I shall not deny that painful negative emotions, when appropriate, possess final value for epistemic or moral reasons. However, I shall argue that this final value is not distinctive; calmer states can possess the very same final value.

⁹ Kauppinen (2017, p. 32).

¹⁰ Ibid, p. 37.

¹¹ Ibid, p. 46.

But before I begin, three points of clarification are in order. First, I do not take my position in this paper to be anti-emotion. In addition to the final value that negative emotions can possess, negative emotions can also be instrumentally valuable, and indeed my own view is that negative emotions play a deeply important role in human psychology. My goal, then, is to resist one way of defending emotions that I take to be problematic.

Second, I do not deny the instrumental role that negative emotions can play in *enabling* us to eventually take on calmer responses to important values. The capacity to calmly recognize and understand injustice, for instance, may causally depend on prior experiences of anger.¹² But the calm understanding of injustice, once it is achieved, can possess the very same epistemic or moral final value that anger can possess, or so I shall argue.

Finally, I focus in this paper on painful negative emotions, as opposed to positive emotions, because I believe that positive emotions, thanks to their pleasantness, *do* have a significant and distinctive final value, a final value that calmer responses cannot possess. Nevertheless, some of the points that I shall make in this paper can extend to positive emotions as well.

With these clarifications in view, it might be wondered what the upshot is of rejecting DFV. As it turns out, rejecting DFV has highly important practical implications for emotion regulation. If painful negative emotions, when appropriate, have a distinctive epistemic or moral final value, then individuals will always have an important *pro tanto* reason to refrain from regulating their fitting, yet painful emotions.¹³ And if this reason is at all weighty, then

¹² See Pettigrove and Tanaka (2014).

¹³ Srinivasan (2018) makes this point about anger.

emotion regulation may be appropriate far less often than we might have thought. By contrast, if DFV is false, then one important obstacle for emotion regulation is removed, and our attitudes towards our emotions may need to change.

I shall begin, in Section 1, by distinguishing emotions from calmer mental states. Then, in Section 2, I shall consider the claim that fitting negative emotions possess a final *epistemic* value that is distinctive. In Section 3, I shall consider the claim that fitting negative emotions possess a final *moral* value that is distinctive. In Section 4, I shall consider the practical implications of rejecting DFV.

1. Emotions and Calm Responses

Philosophers disagree about the nature of emotions.¹⁴ But for the purpose of defending DFV, these differences may not matter. Defenders of DFV need to identify a feature of occurrent negative emotions that 1) sets them apart from calmer mental states, and 2) can explain their distinctive final value, when they are fitting. As the quotations in the introduction section suggest, it is the *affective* or *feeling* component of occurrent negative emotions that seems to play this role. Only when we are appropriately *pained* by injustice, loss, or personal wrongdoing do we properly recognize and appreciate these values, or properly express our concern for them.

¹⁴ There are judgmentalist or quasi-judgmentalist theories (Nussbaum, 2001, Solomon, 2007), perceptual or quasi-perceptual theories of various kinds (de Sousa, 1987, Roberts, 1988, Prinz, 2004, Tappolet, 2005, Doring, 2007), feeling theories (Kriegel, 2014), attitude theories (Deonna and Teroni, 2015), adverbial theories (Berninger, 2016), impure somatic theories (Barlassina and Newen, 2014), and *sui generis* theories (Bahlul, 2015), just to name a few.

But it would be a mistake to claim that calm evaluative judgments are, by contrast, devoid of affect or feeling. Evidence from experimental psychology suggests that affect is ubiquitous.¹⁵ Even walking down the street, we experience slight affective reactions to the various stimuli that we encounter. And calm evaluative judgments, in particular, need not be cold or disinterested, as we shall see.¹⁶ Nevertheless, there is a clear difference between calmly judging that some proposed legislation, say, is *very, very unjust*, and being *outraged* by this fact. So, although calm evaluative judgments are not devoid of affect or feeling, there does seem to be an important affective difference between painful negative emotions and calm judgments. And this is enough to mark off DFV as an important thesis. The defender of DFV can claim, for instance, that only those who are truly outraged by a very unjust piece of legislation fully appreciate how unjust it is; those who correctly, but calmly judge it to be very unjust are missing something.

Now, on one conception of emotions, emotions just *are* a certain kind of evaluative judgment (e.g. Nussbaum, 2001). Strictly speaking, so long as we can distinguish between calmer and more affectively charged judgments, the arguments that I present in this paper are compatible with this view. Nevertheless, I think this view should be rejected. One problem with simple judgmentalist theories is that they have trouble accounting for the phenomenon of

¹⁵ See Zajonc (1980, 1985).

¹⁶ Philosophers sometimes make it seem as though the alternative to an emotional response is a cold, indifferent judgment. For example, Srinivasan (2018, p. 132) contrasts the person capable of feeling anger with “a person who does everything, as it were, by the ethical book ... but who is left entirely cold by injustice, feeling nothing in response to those moral wrongs of which she is perfectly aware.” Similarly, Jollimore (2004, p. 339) asks us to imagine a person, Melissa, “who is incapable of grief and who ... fails to feel much of anything when her best friend, Bob, dies in an accident.” My suggestion here is that such contrasts are misleading. In between painful emotions and cold judgments are warm, evaluative judgments that are not devoid of feeling.

recalcitrance.¹⁷ The judgmentalist would have it, for instance, that occurrent fear of some insect is a judgment that danger is present. But it seems possible for the individual to judge that the insect is *not* dangerous, while still experiencing recalcitrant fear. So one can have an emotion, it seems, without forming the corresponding evaluative judgment.¹⁸ It also seems possible to form an evaluative judgment without experiencing the corresponding emotion. If one calmly judges that a terrible injustice or loss has occurred, this is a far cry from anger or grief. If all of this is right, then painful negative emotions do not only involve more intense affect or feeling than calm evaluative judgments, they are also instances of a different type of mental state.

In summary, then, for the purpose of evaluating DFV, I shall take negative emotions to be distinct from calm evaluative judgments. Emotions are not judgments, but a different type of mental state that characteristically involves strong affect or feeling. Calm evaluative judgments, though not devoid of feeling, fall short of the affective experience that characterizes painful negative emotions. I shall seek to determine whether such painful experiences can provide negative emotions with a distinctive final value, a value that calmer states cannot possess.

¹⁷ For discussion, see Greenspan (1988) and D'Arms and Jacobson (2003).

¹⁸ One could insist here that, in fearing the insect, one *does* form the judgment that the insect is dangerous, and this judgment simply conflicts with one's *considered* judgment that the insect is not dangerous. But this would be to posit a judgment that is highly atypical. The fear-based judgment would be a judgment that the individual does not really endorse and does not at all wish to justify. It seems better to say that the so-called "judgment" is not really a judgment at all, but a different kind of mental state.

2. Epistemic Considerations

A common way of defending the final value of negative emotions relies on the idea that negative emotions can be ways of *recognizing* or *appreciating* important concerns. Lynne McFall, for instance, argues that bitterness can be a way of “bearing witness” to injustice.¹⁹ Macalester Bell makes a similar point about contempt, suggesting that contempt is a way of bearing witness to “deficits in character.”²⁰ Randolph Clarke suggests that the experience of guilt can be “noninstrumentally good ... there is value in the recognition by one who is blameworthy ... that she is so blameworthy.”²¹ Deonna and Teroni claim that if we think that evaluative knowledge is a final good, then we “should extend this claim to the emotional underpinnings”²² of evaluative knowledge. And, as we have seen, Srinivasan, Jollimore, Bok, and Kauppinen defend either anger, grief, or guilt as ways of recognizing or appreciating injustices, losses, and personal wrongdoings. But is the value of emotionally recognizing or appreciating important concerns distinctive – is it a value that calmer states cannot possess?

Let us begin with recognition. Consider the case of Sasha, who is a staunch advocate of gender and racial equality. She is an activist, and she takes her role seriously, traveling across the country and sometimes the world, speaking out against injustice. Sasha has studied the underpinnings of inequality for years, and she knows all about systemic power imbalances, the psychology of bias, and so on. She can speak about all of this quite eloquently, therefore, in good academic fashion. But in her compelling presentations on these issues, her forceful

¹⁹ McFall (1991, p. 155).

²⁰ Bell (2005, p. 88).

²¹ Clarke (2012, p. 155).

²² Deonna and Teroni (2016, p. 171).

commentary never rises to the pitch of *anger*. Even so, it would be odd to claim that Sasha doesn't recognize the injustice of inequality. If anyone recognizes this injustice, Sasha does. Sasha cares deeply about inequality and has devoted herself to the issue. Occurrent anger at the injustice of racial and gender inequality would possess final value, as a recognition of an important aspect of reality. But Sasha's calmer recognition of injustice seems to possess final value too, and for the same reason.

Of course, it could well be that anger is what brought Sasha to this point. Anger may have motivated Sasha to pursue the fight against injustice. It may have motivated her to reflect on the nature of injustice and to study it, so that now she can think and speak about it calmly. It may also be the case that even now, when Sasha encounters a new instance of injustice, she briefly responds with anger or resentment before settling into her usual calm, but determined, approach. In other words, anger may be a stepping stone to a calmer recognition of injustice. And in this way, anger is instrumentally valuable. The question is whether anger has a distinctive *final* value over and above a calmer response.

It might be claimed that there is indeed something missing in Sasha's calmer response, however admirable she may be. If Sasha is not angered by injustice, then she fails to *appreciate* its awfulness. What does such appreciation involve? Well, one aspect of appreciation worth noting is its attentional character. When we appreciate a work of art, we do not typically absorb its beauty or meaning all at once, but rather *attend* to the various features of the piece that make it beautiful or give it meaning. Perhaps something similar happens in the case of emotional appreciation. When we grieve, for instance, we are not only pained by our loss, but our attention is characteristically directed towards the various ways in which the loss affects us.

Consider the case of Fahad, who has just endured the death of his mother. His grief takes him through a prolonged process that slowly directs his attention to the different aspects of his life that must now be different, given the loss. Fahad is initially consumed by the realization, for example, that he will never again have the kinds of basic interaction with his mother that he has always cherished – conversations, walks downtown, Sunday lunch at home. His mind slowly discovers new considerations, new ways in which his life will be affected. In this way, he not only recognizes that he has suffered a loss, but also appreciates the nature and significance of this loss.

But of course, we can also *calmly* attend to the various ways in which an injustice or loss is awful. Sasha can calmly reflect, for example, on the various ways in which the latest injustice is terrible: it perpetuates stereotypes, stigmatizes certain groups, and empowers the wrong set of people. “Yes,” Sasha might think, “that’s what makes this so terrible.” Fahad, similarly, in a calm moment, might reflect on the nature of his loss without being actively pained by it. He might mull over the various considerations he uncovered while grieving. Summarizing, he might think to himself, “I have been affected by this in so many ways.” And eventually, after his grieving period comes to a healthy resolution, it may be that calm reflection becomes Fahad’s default way of recognizing his loss and attending to the various ways in which it has affected him. We have not yet identified, then, the distinctive feature of emotional appreciation. Though emotional appreciation characteristically involves attending to the various ways in which an evaluative object is good or bad, one can attend to such considerations calmly as well.

The missing piece that I have been isolating, of course, is *affect* or *feeling*. The emotional appreciation of negative values involves more than just a characteristic pattern of

attention; it also involves registering negative values in a way that is suffused with feeling. And such affect-laden appreciation might be distinctively valuable. Following Srinivasan, for example, we might compare the appreciation of loss or injustice to the appreciation of beauty or the sublime. It's one thing to calmly *recognize* that beauty exists, but another thing entirely to *appreciate* it. If so, then perhaps painful negative emotions can have a distinctive final value after all. Only by experiencing negative emotions can we fully appreciate important concerns.

In responding to this position, we should begin by noting a deep asymmetry between the appreciation of beauty and the appreciation of loss or injustice. The appreciation of beauty is usually *pleasurable*. In fact, the appreciation of positive values is in general associated with pleasure. Whether one is laughing at a funny joke, celebrating a genuine victory or achievement, or looking forward to an exciting event, the appreciation of positive values is usually enjoyable. Since pleasurable experiences are usually thought to have final value, it's no surprise that the often-pleasurable appreciation of positive values is thought to have final value as well. By contrast, the appreciation of negative values such as loss or injustice is usually painful. This is not always the case, of course, and sometimes it can actually feel *good* to be angry, to let out one's pent-up rage. But I don't think this is what Srinivasan has in mind when she speaks of the appreciation of injustice. Such appreciation, the thought goes, requires one to be *pained*. But if so, then the analogy to the appreciation of positive values is pretty misleading.

So we need another way of investigating this issue. We need to better understand the epistemic import of affect or feeling; we need to better understand the role that such affect or feeling plays in relation to calmer deliberation.

If we look at empirical conceptions of affect, of the feelings associated with emotions, we often find a picture of affect as the mind's rapid evaluation mechanism. Emotional feelings, that is, are often thought to be rapidly-formed and consciously accessible *representations* of value. For instance, Barrett, Mesquita, Ochsner, and Gross (2007) suggest that the "core affect" underlying emotions represents a kind of "knowledge about whether objects or events are helpful or harmful, rewarding or threatening, calling for acceptance or rejection."²³ Similarly, the somatic marker hypothesis proposed by Antonio Damasio suggests that emotional feeling states encode information about the goodness or badness of various choices, guiding decision-making.²⁴ And dual-process theories assimilate emotional responses into so-called "Type 1" responses, which are typically rapid, effortless, and autonomous, in contrast with the slower, more effortful, and consciously directed Type 2 processes associated with higher cognition.²⁵ Affectively charged emotional processes are understood as initial evaluative assessments that can then be challenged or modified by higher-order conceptual processes.²⁶ Affect, then, is not conceived as epistemically profound, but as providing a rapid evaluative assessment that can then be interrogated by slower, more costly cognition.

This conception of affect as representational in nature extends to the philosophical literature. Peter Carruthers, for instance, has recently argued that the *valence* of affective

²³ Barrett, Mesquita, Ochsner, and Gross (2007, p. 377). See also Russell (2003, p. 148): "feeling is an assessment of one's current condition."

²⁴ Damasio (1994).

²⁵ See, for instance, Bargh and Chartrand (1999), Haidt (2001), and Cushman, Young, and Greene (2010).

²⁶ Affectively charged processes are thought to be *initial* assessments in the sense that they are the first assessments of novel situations or events. But these assessments are typically made on the basis of information learned from prior experiences, and so are not *initial* assessments in any deeper sense.

states is a nonconceptual representation of value.²⁷ And Peter Railton points out that “affect has many of the features necessary for the representation of value: It varies in degree and intensity; it can take a wide range of intentional objects (e.g., one can be afraid of a situation, person, action, or risk); it possesses positive or negative valence ...”²⁸ Now, emotional feelings do not need to be construed as particularly *crude* representations of value. On the contrary, Railton argues that affective representations are generated by neural systems that are capable of “tracking statistical dependencies”²⁹ and of weighing “costs, benefits, and risks.”³⁰ And these systems are sensitive to new information, constantly attuning us to our environments.³¹

Importantly, such affective processing seems to be ubiquitous. And so, we can find examples of the transition from feeling-based representations to conceptual representations even during calm, philosophical deliberation. As we reflect on some particular philosophical issue, our affective system is active, and so we may *feel* that certain lines of inquiry are promising, others doubtful. And this is often a very good thing. Affective representations may summarize a wider range of information than is consciously accessible at a specific moment in time. We may, for instance, feel that there’s something wrong with our current line of reasoning without being able to say just what it is. Something about one of the premises just doesn’t seem right. But we can unravel this concern and put our uneasiness about the issue

²⁷ Carruthers (2017).

²⁸ Railton (2017, p. 339).

²⁹ Railton (2014, p. 832).

³⁰ *Ibid.*

³¹ *Ibid.*, p. 832-839. According to Railton, “the full statistical models generating these feelings would be much too complex to represent in consciousness, with its narrow ‘band- width’ and limited working memory. But conscious representations of these models might not be necessary if conscious *feeling* can present the relevant information in a summary way.” (p. 340, italics in the original).

into words. We can think the problem through. And in this way, by putting our feelings into words, we can shift from an inchoate *sense* or *feeling* that something is amiss to a conceptual representation of the matter at hand. In doing so, we may reach a conclusion that either confirms our initial feeling that something was wrong (“yes, that premise is false, on reflection”), or that challenges somehow our initial feeling-based representation (“no, that premise only seemed dubious because I misinterpreted it.”).

Consider now the case of *emotional* feelings. Emotions provide us with rapidly-formed representations of value that can immediately guide our actions. Although sophisticated conceptual representations of the values or concerns in question can be worked out later, feeling-based representations can come online immediately, allowing us to respond rapidly and without reflection. Anger, for instance, can represent moral violations, and so can rapidly mobilize individuals to confront injustice. If the government engages in some egregious act, the streets can be filled with protestors long before an academic analysis of what the government did can be developed. Now, even when urgent action is not clearly necessary, intense emotional feelings provide us with representations of value that can dominate our consciousness and so guide our thoughts and actions in a certain direction. Grief is an orchestrated, prudentially important process that represents the horribleness of a personal loss, and the all-consuming nature of this representation can direct the individual to reflect on the loss, to unravel its implications, and so to come to terms with and adjust to the loss. In this way, grief can prompt the individual to engage in a process that produces a conceptual representation, an understanding, of the loss and its significance.

The feeling-based representations we find in anger and grief, then, are rapidly formed, conscious representations of value that fill the role that more sophisticated conceptual representations can play later on. Sasha, for instance, possesses a sophisticated understanding of racial and gender inequality. Anger may revitalize her motivation to keep fighting against this injustice, and so it may remain instrumentally valuable for her. But Sasha does not lose anything of final epistemic value by regulating her anger and deploying her more sophisticated conceptual representation of injustice instead. Similarly, when Fahad, in a cooler moment, reflects calmly on the nature of his loss, his understanding of the loss need not be any less profound than when he was grieving. In fact, his calmer reflection may contain a specificity that his grief did not; he may be able to put into words just how this loss has affected him.

In light of these considerations, it seems that the *appreciation* of negative values through strong feelings is the mind's way of initially, rapidly, and immediately representing such values and mobilizing the individual to respond. But the representations underlying emotional feelings can then be challenged, modified, and incorporated into higher-order conceptual representations of value. Individuals can then access these representations of value without reigniting intense feelings. If this is right, then intense feelings are not necessary for representing the severity of an injustice or a loss. Emotional appreciation of important concerns may very well have final value, since such appreciation is a way of recognizing important concerns. But calmer reflection on important concerns, involving higher-order conceptual representations, can possess the very same final epistemic value.

Now, as I have already mentioned, calm evaluative judgments are not devoid of affect or feeling, which seems to be ubiquitous; our affective system is constantly active, constantly

evaluating. So it is not surprising that calm *evaluative* judgments involve some affect or feeling, and perhaps this connection is even necessary – perhaps an evaluative concept wouldn't *be* evaluative if it didn't tend to activate the affective system in a certain way. Still, though calm evaluative judgments involve affect, they can, as I have said, represent highly important concerns or values without *intense feeling*. We possess a higher-order representational system that allows us to represent, using evaluative concepts, the very same evaluative features that our intense feelings represent. We can represent what our *outrage* is about by speaking of a "terrible injustice," and we can represent what our *grief* is about by speaking of a "horrible loss."

It would be extremely surprising if intense feelings could represent certain aspects of evaluative reality that we could not represent in language. Given the role that intense feelings seem to play in our psychology, as rapidly formed evaluations that can then be challenged or modified using our conceptual system, we should not expect intense feelings to be capable of representing anything that higher-order representations cannot.³²

Why, then, have some philosophers thought that emotional feelings are necessary for truly recognizing and understanding important concerns or values? I think there are a few forces at work here. First, given the way that our psychology is organized, affective registration of important concerns or values is often the only way to *initially* grasp the magnitude of these

³² It might be objected that emotional feelings represent the magnitude of values more precisely than coarse-grained conceptual representations. Even if this is generally true, the calm judgment as a whole must at some level represent the fine-grained magnitude of the value, because even calm people can make fine-grained decisions between options that are conceptualized as "very bad." So the emotional representation does not possess any distinctive epistemic final value that a calmer mental state, as a whole, cannot possess.

concerns or values. Perhaps an expert like Sasha can grasp the enormity of some injustice calmly, even initially. But Sasha is statistically rare. And so, when we are asked to imagine a person who is not angered by injustice, we typically imagine a person who doesn't fully register the injustice, or who does not sufficiently *care*. We have a paradigm, then, of the calm responder, and this paradigm seems to be powerful – I think it infects even our perception of Sasha. We are tempted to think that Sasha, insofar as she is calm, must have grown cold to injustice, that she must no longer *fully* appreciate just how horrible certain injustices are. And we are tempted to think this even after we are told that Sasha has a deep understanding of injustice and is highly motivated to resist it.³³ In other words, we are dealing here with a powerful, recalcitrant intuition that nevertheless cannot be justified in cases like Sasha's. The strength of this intuition suggests that we perceive lack of emotional responsiveness to be *highly informative*, highly indicative of a defect in understanding or character. But although this intuition may prove robust and reliable across a range of cases, it is important not to draw the wrong lesson from the strength of this intuition. It is true that people who fail to respond emotionally to important concerns, at least initially, usually do not grasp the importance of the concern in question. But it does not follow that intense emotional feelings provide some kind of unique epistemic access to evaluative reality. On the contrary, I have argued that they do not.

³³ This suggests that the phenomenon we are dealing with here may have some relationship to the moral dumbfounding effect in moral psychology (see Haidt, 2001). Moral dumbfounding occurs when individuals form judgments on the basis of strongly held intuitions, but then are unable to justify their judgments in the face of contradictory evidence. In such cases, individuals find themselves “morally dumbfounded.” They may “stutter, laugh, and express surprise at their inability to find supporting reasons, yet they [do] not change their initial judgments” (Ibid, p. 817). More charitably, however, some have argued that the intuitions underlying moral dumbfounding may be picking up on important information that is in fact relevant to the situation at hand (see, for instance Railton, 2014, and Uhlmann, Pizarro, and Diermeir, 2015). As I go on to explain, I think it is indeed important to acknowledge the information that is carried by our powerful intuitions; but it is also important to take care not to draw the wrong lesson from these intuitions.

Once we affectively register an important concern, we are often able to regulate our emotions and calmly reflect on and understand the concern in question.

Second, I think it is easy to mistake two different things: the difficulty of describing, in language, what our feelings of outrage or grief are *like*, and the difficulty of describing, in language, what our feelings of outrage or grief *represent*. Jollimore, for instance, suggests that the person who does not experience appropriate grief is missing something, but Jollimore is willing to grant that what she is missing cannot be expressed in a proposition.³⁴ Jollimore compares such a person to the person who has never seen the color red, and so does not know what it looks like. But we must be careful to distinguish what the experience of the color red is *like* from what this experience *represents*. Although it is difficult to put into words what the experience of the color red is like, it is very easy to put into words what this experience represents: namely, a particular shade of red (e.g. #FF0000). Similarly, we can of course agree with Jollimore that the person who has never felt grief may not know what it is *like* to grieve. But this is a far cry from the claim that what grief represents cannot be expressed in a proposition. Though we can disagree on the exact formulation, we can agree that grief is (roughly) an experience that represents a significant, personal loss as truly horrible.³⁵ As I have been arguing in this section, we can *easily* represent in language what our feelings are about, thanks to the way in which our psychology is organized.

³⁴ Jollimore (2004, p. 341-342).

³⁵ Michael Cholbi (2017) suggests that grief is about the loss of a personal relationship. If so, then we can say that grief represents the loss of a personal relationship as truly horrible.

If all of this is right, then painful negative emotions do not possess a distinctive, epistemic final value.³⁶ Still, it might be felt that there is nevertheless something distinctively valuable about experiencing appropriate negative emotions. Perhaps there is a distinctive *moral* final value to being pained by important concerns. In the next section, I shall investigate this possibility.

3. Moral Considerations

There are a number of ways in which one might defend the final moral value of negative emotions. A straightforward approach would be to argue that we have a pro tanto duty to experience certain appropriate negative emotions. In the case of grief, for example, perhaps we have a pro tanto duty to be pained by our significant losses; perhaps Fahad has a pro tanto moral duty to grieve the death of his mother.

It's worth noting that when philosophers write about grief, they tend not take this route, opting instead for an epistemic defense of grief. According to Dan Moller, for instance, one reason to regret the fact that individuals recover from grief relatively quickly is that this "renders us unable to take in and register fully the significance of our losses."³⁷ Cholbi argues that "grief is a particularly fruitful opportunity for an important human good: self-knowledge."³⁸ And Jollimore, quoted already, suggests that grief "is a way of understanding

³⁶ My argument in this section depends on the view that emotional feelings represent value. Could the defender of DFV push back by denying this premise? I think not. On the contrary, endorsing this premise, at least for the sake of argument, is the charitable thing to do on behalf of the defender of DFV. After all, if emotional feelings do not represent – if they are mere qualia – then it is hard to see how they could possess final epistemic value.

³⁷ Moller (2007, p. 310).

³⁸ Cholbi (2017, p. 101).

what has happened in one's life, and more broadly, of appreciating the nature of the world in which our lives take place."³⁹ Now, it could be that we have a pro tanto moral duty to attain certain epistemic achievements. Perhaps we have a pro tanto duty to recognize and understand the extent of our personal losses. But as I have argued in the previous section, grief would not be the only mental state that fulfills this duty. We can calmly recognize and understand the extent of our losses too.

But perhaps it is possible to defend the final moral value of grief independently of epistemic considerations. Perhaps there is final value in *honoring* those we have lost, and perhaps grief is the only way, or the best way, to honor the deceased.⁴⁰ One might argue that the most respectful way to honor a loss is to be pained by it. The idea here is not that we have a pro tanto duty to suffer *simpliciter* – it wouldn't do to downregulate our grief and then self-flagellate. Rather, the thought is that we have a pro tanto duty to be pained by the loss itself – that is, to grieve.⁴¹

But this account requires further explanation. After all, there are many ways of honoring those we have lost: we can memorialize them, we can fulfil their last wishes, and we can calmly reflect on what they meant to us. Once again, we return to the question of why an intense *affective* response is especially valuable. *Why* should suffering be the most respectful way to honor a loss? What underlies this intuition? Perhaps it is this: our suffering shows that we cared

³⁹ Jollimore (2004, p. 345).

⁴⁰ This characterization may make it seem as though grief is instrumentally, rather than finally, valuable. But I don't think that's right. The idea here is that to grieve *is* to honor, not that grieving brings about honoring as a consequence. Similarly, joy *is* a form of happiness, not an instrument that brings happiness about as a consequence.

⁴¹ For related discussion, see Bok (1998) and Clarke (2012).

about the person we lost, that they were important to us; moreover, our suffering can manifest our understanding of these facts.

Unfortunately for the defender of DFV, I do not think it is plausible that the *pain* of grief is necessary for any of this, or more valuable than calmer alternatives. Recall that Fahad, in a cool moment, can calmly reflect on the significance of his loss. And suppose that he is highly motivated to do so. Calmly, he reflects on what a profound loss he has suffered, about all of the various ways in which it will affect him. He would do anything to bring his mother back, he thinks. His mind naturally turns to how he might honor his mother. He begins to reminisce on happy memories that they shared. He resolves to follow up on some of the projects his mother left behind. It would have meant a lot to her. And so on. In reflecting in these ways, and being motivated to do so, Fahad shows how much he cares, how important his mother was to him, and how he understands the depth of his loss. Of course, the pain of grief can show these things too; and so the pain of grief is also a way of honoring his mother. But it does not seem to be the only or the best way.

Now, none of this is to deny the claim that in human beings, a failure to grieve *indicates* a lack of caring. We are wired to grieve upon significant losses, and so if we don't grieve, that means something. This may well explain, at least partially, the strong intuition that occurrent grief is the only way, or the best way, to honor a loss or show that we cared. When we are asked to imagine a calm response to a significant personal loss, we are inclined to imagine the person who doesn't care, or the person who hasn't yet registered fully the enormity of the loss. We do not typically imagine someone like Fahad, who does understand the magnitude of his loss, and who, in a cool moment, is able to show that he cares by calmly reflecting on the loss

and resolving to honor it in various ways. Yet Fahad's case is real and important. And so we can accept the conclusion that a failure to grieve may indicate a lack of caring; but it does not follow that the pain of grief is the only way to honor a loss or show that we care. On the contrary, it seems that calmer ways of honoring a loss can possess the very same final moral value.

Moving on from grief now, perhaps a more forceful argument can be made in defense of appropriate *guilt*. After all, we might think that those who are blameworthy for committing some wrong *deserve* to suffer on account of what they have done. And we might also think that, in contrast to the case of grief, no calmer substitute would possess as much final moral value. If the wrongdoer does not suffer, they haven't *paid* for what they've done, in some sense. Appropriate guilt possesses a distinctive final value.

Once again, philosophers tend not to defend guilt in this way, relying instead on the value of correctly recognizing one's culpability. As quoted earlier, Clarke defends the final value of guilt by suggesting that "there is value in the recognition by one who is blameworthy ... that she is so blameworthy."⁴² Similarly, A.B. Carlsson writes: "When you feel guilt, you suffer in recognizing what you have done. I believe that this captures the specific kind of suffering that we take blameworthy agents to deserve. It is not suffering in general, but rather the specific pain of acknowledging your own objectionable quality of will."⁴³ And, Bok argues that the suffering associated with guilt does not require "its own separate justification,"⁴⁴ since such suffering is caused by "the recognition that one has done something wrong."⁴⁵ But as I have

⁴² Clarke (2012, p. 155).

⁴³ Carlsson (2017, p. 105).

⁴⁴ Bok (1998, p. 170).

⁴⁵ *Ibid*, p. 168.

argued in the previous section, though guilt may have the final value of recognition, this final value is not distinctive, since calmer reflection can possess this value too.

Let us return, then, to the idea that appropriate guilt possesses a distinctive final value because it is the only proper way of paying for what one has done. Once again, I suggest that this position is implausible. The pain of guilt is not the only way, or the best way, to pay for one's wrongful actions.

Consider the case of Lydia, for example. Lydia had promised her friend, Misa, that she would come and perform at her big, annual charity party. This was an important role, as Misa hoped that Lydia's talent would help to stimulate donations. Nevertheless, on the day of the charity party, the obligation slipped Lydia's mind. Lydia lay down for a nap and slept through her scheduled performance. She woke up to dozens of text messages and missed calls: "Where are you?!?!" Immediately, Lydia felt enveloped with guilt. She recognized her utter carelessness – this is something she should have safeguarded in some way, with reminders, plans, and so on. But she hadn't. She had simply assumed that everything would be okay. Lydia let out a groan and covered her face. But at that moment, her regulatory habits kicked in. She took a deep breath and told herself that there was no great mystery about what she had to do. And there was no need to wallow in the pain of guilt: "I'm a human being, and human beings make mistakes. And I made a big mistake, one that I need to make up for, pay for, and learn from." Lydia immediately called her friend Misa, ready to apologize sincerely, and to determine whether it was too late to perform.

Since that fateful day, Lydia carried into action her plan for atonement. First, she changed her scheduling and reminder protocol, creating a system that would surely prevent similar occurrences in the future. Then, she not only helped Misa to make up for the missed donations opportunity by performing at a different event, she also volunteered for a larger role in Misa's charity organization. This is something that Lydia didn't particularly look forward to doing – such things usually bore her tremendously. But Lydia recognized that it wasn't enough to simply undo the damage that she had done. To truly make it up to Misa, Lydia would have to show that she was willing to go further, to truly pay for her mistake. Now, throughout this process, Lydia did from time to time feel pangs of guilt, as she experienced flashbacks of her careless mistake. But when this happened, Lydia immediately regulated her guilt, reminding herself that she was already perfectly aware of what she had done, and that she was already working overtime to pay for it – again, no need to wallow in the pain of guilt.

Now it might be objected that if Lydia does not suffer during this process of atonement, then she does not *fully* get what she deserves. But this position seems highly implausible to me. Even if Lydia remains calm throughout this process, she is driven by a sincere desire to take responsibility, to correct her character, and to make things right with Misa. Lydia is willing to work on this as long as it takes. She is willing to take time away from her own pursuits, her own plans, and the activities that she enjoys in order to make things right. Eventually, it will seem odd to insist that she still hasn't paid the price for her mistake. Now, in saying this, I do not have to deny that guilt *would* have final moral value, that guilt *would* be a way of "paying the price." I simply claim that there are other ways of responding that possess the same final moral value, because they are also ways for Lydia to pay for what she has done. If this is right, then

appropriate guilt does not have a distinctive final moral value. Individuals can calmly recognize, understand, and pay for their wrongful actions.⁴⁶

Let us consider one last way of defending the final moral value of negative emotions. Some philosophers suggest that emotions possess final value when they manifest our *virtuous concerns* or our *respect for others*. These emotional manifestations of concern and respect, the thought goes, are valuable even when they do not result in any further good. For instance, Deonna and Teroni point out that emotional responses manifest our virtues.⁴⁷ Virtues, Deonna and Teroni suggest, are valuable as ends in themselves, and since “emotions count among virtues’ most fundamental manifestations ... it is difficult to dispute the conclusion that emotions can be finally valuable.”⁴⁸ Similarly, Charlie Kurth argues that emotions such as anger and anxiety can express virtuous moral concerns: “anxiety and anger ... are valuable because of what they say about one’s emotional attunement. Importantly, the value here isn’t (merely) instrumental in nature ... to have the virtue of moral concern is to be a person who feels angry when what they care about is threatened and who feels anxious in the face of a difficult or

⁴⁶ Might there be some crime so heinous that a calm response would never suffice? This is a difficult question, because it’s hard to wrap one’s mind around the relevant thought experiment: an individual commits a heinous crime, but eventually comes to fully understand the horrific damage he has caused and resolves to spend the rest of his life making up for it. But this individual regulates his guilt, much like Lydia. Would it be *better* for this individual to also be in pain? Well, it’s hard in the first place to imagine this individual, who achieves such a dramatic change of heart, but insofar as I am able to do so, I do not feel that a response involving more pain would have greater final value. This individual, however terrible he once was, now (by hypothesis) fully understands the wrong he has committed and is fully motivated to atone. So further pain will not deepen his understanding or motivation any further. He can pay for his past ways calmly, without the pain of guilt.

All this being said, it is not my position that a calm response, possessing as much final value as a negative emotional response, is always *available*. Perhaps the individual who commits a heinous crime will simply not live long enough to make up for it. But I do not deny this. Instead, I only seek to reject DFV. And rejecting DFV is compatible with granting that a sufficiently valuable calm response may not always be available, for contingent reasons.

⁴⁷ Deonna and Teroni (2013, p. 132).

⁴⁸ *Ibid.*

novel choice.”⁴⁹ In a similar spirit, Christopher Evan Franklin argues that blame can manifest or express our moral valuations: “we should be disposed to experience blame toward the agent who has devalued an object of moral value ... in failing to blame the murderer, we (all things being equal) fail to value the victim.”⁵⁰ Relatedly, Antti Kauppinen, quoted in the introduction, argues that anger, as a reactive attitude, can express our respect for the personhood of others, in contrast with a Strawsonian objective attitude: “One possible response is to treat you as a nuisance to be managed ... But taking such an objective attitude, in Strawson’s (1962) terms, is not consistent with respecting you. Respect requires instead that I’m willing to hold you accountable ... I must show my anger and resentment.”⁵¹

There does indeed seem to be something valuable about expressing our virtuous concerns and our respect for others. But I do not think that painful negative emotions are necessary for this. In fact, we have already encountered this point in our discussion of Fahad and his grief. We noted that Fahad is able to express his virtuous concern for the loss of his mother even when he is not actively experiencing the pain of grief. Fahad, in a calm moment, finds himself highly motivated to reflect on his loss, to cherish his mother’s memory, and to think about how he might honor his mother by pursuing projects that were important to her. In being so motivated, Fahad expresses his virtuous concern for the loss of his mother. In general, I think that we can express our concerns quite well through our motivations and actions. It may be a contingent fact that if, as human beings, we do care, we will also sometimes experience

⁴⁹ Kurth (2018, p. 101).

⁵⁰ Franklin (2013, p. 217). See also Scanlon (2013): “for us, as third parties, simply to ignore the attitudes toward the victim (and perhaps others) that are indicated by the agent’s actions would show a lack of regard on our part for the moral standing of these people” (p. 106).

⁵¹ Kauppinen (2017, p. 37).

negative emotions. But again, it does not follow that experiencing these painful emotions is the only way, or the best way, to express our virtuous concerns. For further examples, recall the cases of Sasha and Lydia. Both of these individuals express their virtuous concerns through their motivations and actions. Sasha is motivated to fight injustice; Lydia is motivated to make up for her mistake. If this is right, then although experiencing painful negative emotions is a way of expressing our virtuous concerns, the final value of this expression is not distinctive.

What about Kauppinen's idea that anger, in contrast with a Strawsonian objective attitude, is a way of expressing our respect for the personhood of others? This point should only concern us if we think that downregulating anger means taking up a Strawsonian objective attitude. Of course, we sometimes *do* regulate our anger by shifting into an objective attitude. But this is not necessary; we can downregulate our anger and still fully respect the agency of others. Suppose that Juan and his partner Q are having a fight. Juan left a mess in the kitchen *again*, and Q, who is less tolerant of such things, wants him to clean it up. But Juan dismisses her: "What does it matter? I'll do it later." Q starts to get really upset. It's not just about the kitchen; it's about the fact that Juan just doesn't seem to be responsive to her concerns. *Even if* her concern is trivial (which it isn't!), Juan should care that *she* cares. But Q knows that Juan will just leave the room if she expresses anger, so she downregulates. But her next step isn't to take the objective attitude, treating Juan as some object to manipulate. Instead, after taking a deep breath, Q tries to calmly reason with Juan: "Look, I know that this kind of thing isn't important to you. Just like how getting to the movies on time to see the ads isn't important to me. But see, I make sure we do get there on time, because I know *you* care. That's all I'm saying." Far from treating Juan as an object to be manipulated, Q is respecting Juan as an agent – she

reasons with him, appealing to his sense of fairness. So although her anger had final value in that it respected Juan's rational agency, this final value was not distinctive; a calmer approach also respected Juan's rational agency, and so possessed the same final value.

In summary, it is difficult to defend the idea that painful negative emotions possess a distinctive final value on moral grounds. I therefore conclude that DFV is false. Painful negative emotions may possess final value for epistemic or moral reasons, but this value is not distinctive; calmer mental states can possess the very same final value.

4. Practical Implications

Rejecting DFV has important practical implications for emotion regulation. If DFV is true, then individuals always have an important pro tanto reason to refrain from regulating their fitting negative emotions. If this reason is at all weighty, then emotion regulation may be appropriate far less often than we might have thought.

Srinivasan, writing on anger, makes this point very well. She points out that although fitting anger may be counterproductive, it may still be the case that we ought to experience this anger, because it is *fitting* or *apt*:

On such occasions, I want to suggest, reasons of prudence and reasons of aptness come apart, generating a substantive normative conflict ... the counterproductivity critic faces the burden of explaining why, in such conflicts, reasons of prudence trump reasons of aptness; until this burden is met, there is no obvious inference to be made from the counterproductivity of one's anger to an all-things-considered prohibition on one's getting angry.⁵²

⁵² Srinivasan (2018, p. 127). Srinivasan uses the terms 'fitting' and 'apt' interchangeably.

An upshot of rejecting DFV is that we can typically avoid the substantive normative conflict that Srinivasan identifies here. If the final epistemic or moral value of fitting anger is a value that can be possessed by calmer mental states, then individuals can downregulate their counterproductive anger without giving up this final value. They can downregulate their anger and calmly reflect on injustice instead. And so, by rejecting DFV, we remove a major normative obstacle for emotion regulation. So long as we regulate our emotions in a way that allows us to continue attending to important concerns, we can regulate without losing anything of epistemic or moral importance (a corollary, then, is that we may well lose something of epistemic or moral importance if we regulate our emotions by trying to ignore or minimize significant concerns).

The practical implication just mentioned applies not only to the downregulation of emotions, but also to the upregulation of emotions. Suppose that Sasha, for instance, is calmly reflecting on the nature of some terrible injustice. If DFV is true, then Sasha has a possibly weighty reason to upregulate her anger in this case, since anger would possess a distinctive epistemic or moral final value. But as I have argued in this paper, this position is mistaken. Sasha has nothing of epistemic or moral final value to gain by upregulating her anger.

Let us think now about grief for a moment. Dan Moller cites evidence indicating that individuals adapt to the loss of loved ones in a remarkably short period of time. One reason to regret this ability to move quickly beyond grief, Moller argues, is that this “renders us unable to take in and register fully the significance of our losses ... our reaction to [these] losses is utterly

incommensurate with their value, especially after the first month or two have passed.”⁵³ Moller may well be on to something here, but I think he would be mistaken to attribute a distinctive final value to grief. First, our capacity to rapidly adapt to significant losses need not, and I believe does not, prevent us from fully understanding the magnitude of our losses. We can fully understand that a loss was immense, and that it will reconfigure our lives and affect us in many ways. And we can reflect on all of this calmly, once we have grasped it, without suffering the pain of grief. Now, Moller may well be right that there is nevertheless something to regret in our capacity to move on from a loss relatively quickly, but again, I do not think that this is best understood as a point about grief. Perhaps we have a moral reason not to move on – not to find a new partner, not to stop thinking about our loved one, and so on. But this need not be understood as a reason to keep grieving. If we do have a reason not to move on, I think this is a reason to keep the loved one in our minds, to reflect on our loss, and to honor it. But none of this requires the pain of grief. There is a calm way of relating to our loss that would not be “utterly incommensurate” with its value.

Lastly, consider Lydia’s guilt. Lydia fully recognizes that she has made a mistake, and she is highly motivated to make up for it. But she downregulates her pangs of guilt whenever they arise. If the arguments I have presented in this paper are correct, then there is nothing objectionable about what Lydia is doing. More generally, we can regulate our emotions (our shame, our guilt, our sadness, etc.) without losing anything of final epistemic or moral value, so long as we recognize and understand important concerns and correctly respond to them.

⁵³ Moller (2007, p. 310).

Conclusion

In this paper, I have argued that DFV is false. Painful negative emotions do not have a distinctive final value, for epistemic or moral reasons, that calmer mental states cannot possess.

It is important to acknowledge, however, that the arguments I have presented in this paper only go so far. In particular, my position leaves it open that there may be non-epistemic, non-moral ways in which painful negative emotions might possess a distinctive final value. By way of concluding this discussion, I would like to gesture at some of these possibilities and their implications for emotion regulation.

First, as I have noted already, I do not deny that it can sometimes feel good to experience negative emotions, for various reasons. It can feel good to let out pent-up emotions; it can feel good to express one's feelings to someone who cares; it can feel good to rage against the system that *doesn't* care, to air one's grievances. And giving in to one's emotions can feel better than striving unsuccessfully to regulate them. This is all true, as far as it goes. But in this paper, I have been arguing that *painful* negative emotions do not possess a distinctive epistemic or moral final value. It is compatible with this point to concede that, when negative emotions feel good, they possess final prudential value on account of their pleasantness. And whenever this is the case, we may well have a compelling reason not to regulate our emotions.

Relatedly, painful *aesthetic* experiences may possess final prudential value. We pay good money to experience works of art which fill us with horror and sadness, for instance, a

puzzling phenomenon known in the literature as the “Paradox of Negative Emotions” or the “Paradox of Tragedy.”⁵⁴ Why do we desire such experiences? Prominent theories hold that painful emotional responses to art are either *converted* into a pleasant experience by some accompanying positive affect (e.g. Hume, *Of Tragedy*), or are *compensated* for by some positive aesthetic experience (e.g. Carroll, 1990). On these accounts, we do not desire to be pained by art as such. An important dissenting view, however, is Aaron Smuts’ Rich Experience Theory, according to which we really do desire painful aesthetic experiences for their own sake, even when these are not pleasant. Perhaps we have an intellectual interest, as Smuts suggests, in experiencing certain negative emotions that we could not experience in real life without major consequences (“We can feel fear without risking our lives, pity without seeing our loved ones suffer, thrills without risking going to jail”⁵⁵). I do not wish to rule out this kind of view, but I do wish to emphasize that its practical significance is limited. In the ordinary course of our lives, outside of our interactions with art, the fact that painful experiences carry a certain aesthetic interest does not seem to be a very weighty consideration in deciding whether we should regulate our emotions. By contrast, were painful emotions to possess a distinctive epistemic or moral final value, the practical import of this fact would be far more significant.

Finally, it might be argued that there is something fundamentally human about affectively registering important concerns. There would be a kind of loss, the thought might go, if any one of us were to transform herself into a creature that could only calmly recognize negative values. Perhaps such a person would not lose anything of epistemic or moral value.

⁵⁴ For a review, see Smuts (2009).

⁵⁵ Smuts (2007).

But she would distance herself from the rest of humanity and leave behind a certain way of being in touch with evaluative reality. Perhaps this is so. But what implications does this have for creatures like us, humans who have not yet transformed themselves into a different kind of being? The practical implications, I think, are few. It is implausible that on any particular occasion, our humanity is a reason not to regulate our emotions. After all, emotion regulation too is a human capacity. If anything, then, we reaffirm our humanity by experiencing emotions that we then regulate. But perhaps the value of our humanity means that the ideal for which we should strive is not permanent equanimity, but something more down to earth. Even if we should often be regulating our painful emotions, perhaps there is something inherently good about the fact that we do feel them, before we regulate. I do not know how far this thought goes, but I certainly do not take myself to have ruled it out.⁵⁶

Whatever one makes of the possibilities just discussed, the takeaway is this. The typical way in which philosophers have defended the distinctive final value of painful emotions is problematic, and so too are the practical implications that are meant to follow. Although there may be other ways of defending the final value of painful negative emotions, the practical implications of these other defenses seem far more limited. The upshot, then, is that an important normative obstacle for emotion regulation is removed. In the scheme of our ordinary human lives, the final value of painful negative emotions should not typically prevent us from seeking a calmer state.

⁵⁶ My thanks to Sarah Buss for helpful discussion on these points.

Conclusion

At last, we are in a position to say something general about emotion regulation and the role that it should play in our lives. I shall begin by summarizing some of the implications from the preceding chapters. As we shall see, the chapters in this dissertation help to clarify some important issues surrounding emotional assessment and emotion regulation, and they pave the way for a regulation-friendly position. But it is Chapter 2, in particular, that lays the foundation for a practical, normative model of emotion regulation.

Let us begin by summarizing some implications. Chapters 1 and 3, when considered together, can help to clarify our understanding of the relationship between fittingness, final value, and emotion regulation.

Chapter 1 explains what it takes for an emotion to be fitting, and so helps illuminate what the value of emotional fittingness might be. According to Chapter 1, an emotion is fitting when it not only correctly represents some value, but also correctly mobilizes the individual to respond to the value in question. The value of emotional fittingness, then, can be understood as the value of correct representation and correct mobilization.

Consider correct representation first. We know from Chapter 3 that calm evaluative judgments can possess the very same final epistemic value that emotional representations can possess. This is because calm evaluative judgments can have the same representational content that emotional representations can have. Put in terms of fittingness, we can say that calm

evaluative judgments can be just as fitting as emotional representations of value – and so just as valuable in that respect.

Consider next the non-representing components of the emotional response. These components can themselves be fitting, according to Chapter 1, in virtue of their correctly heeding the call of some evaluative concern (e.g. a call for avoidance, a call for resistance, etc.). There may well be some final value that goes along with fittingness of this sort – suppose, for the sake of argument, that there is. Even if so, Chapter 1 suggests that non-emotional alternatives are available. For *action-responses*, just like the non-representing components of emotions, can be fitting responses to the target value; they too can correctly answer the call of the evaluative concern in question.

Taken together, this analysis confirms that we can regulate our fitting emotions without losing out on any fittingness-related final value, simply by engaging in an alternative fitting response. Instead of feeling anxious, we can calmly recognize and prepare for threats. Instead of feeling angry, we can calmly recognize and confront injustice. And instead of feeling guilty, we can calmly recognize and atone for our mistakes. This complements nicely the conclusion of Chapter 3: we can downregulate our fitting, painful emotions without losing any epistemic, moral, or fittingness-related final value.¹ So long as we continue to recognize and correctly attend to important concerns, we need not lose anything of final value when we regulate our painful emotions.

¹ Remember that this statement must be restricted to painful emotions, because pleasant emotions do have a distinctive final value that we can lose by regulating.

So much, then, for final value. What about instrumental value – in particular, the instrumental epistemic value of emotions? According to Chapter 2, emotions can enhance our evaluative understanding by mobilizing directed reflection. But, as I argued in Chapter 2, in order to achieve full evaluative understanding, we must ultimately regulate our emotions in a particular, organized way. So even the instrumental epistemic value of emotions is compatible with a regulation-friendly position.

~~~

With these implications in view, we might begin to speculate about what a positive normative model of emotion regulation might look like. Well, it is Chapter 2, in my view, that lays the foundation for a particularly fruitful way of organizing our thinking about when and how we should regulate our emotions. And so, by way of conclusion, I would like to gesture towards a practical, normative model for emotion regulation that takes Chapter 2 as its starting point. Although I do not have the space here to fully explore this model, I hope that the brief sketch that I shall provide will be helpful in highlighting an important future direction for this project.

We can think of emotion regulation as playing two roles. First, as explained at length in Chapter 2, emotion regulation can help us to enhance our understanding of an evaluative situation. Once an understanding of the situation is achieved, we are in a good position to form a plan regarding what to do. Having formed this plan, emotion regulation can play another role: it can help us to adhere to this plan and execute it successfully – or, in other words, to act in accordance with our understanding.

This conception of emotion regulation as having two roles mirrors the dual instrumental role that emotions have in our lives: emotions 1) help us to understand the nature of the



evaluative situations that we find ourselves in and 2) mobilize us to confront these situations. Emotion regulation simply helps us to steer our responses in the right direction. Emotion regulation can help us to move beyond our initial emotional assessment of some evaluative situation. And emotion regulation can help us to better mobilize to respond to the situation at hand.

As I argued in Chapter 2, both engaging and disengaging forms of emotion regulation have a role to play in enhancing evaluative understanding. But both types of regulation also have a role to play in helping us to execute our informed plans. These plans tell us which values we should pursue, which concerns we should address, and when all of this should happen. Disengaging forms of regulation can help us to let go of concerns that we have decided we should not deal with, at least for the moment. And engaged reappraisals can help to remind us of what we already understand, whenever our emotions seem to lead us astray. Finally, once we have disengaged from the concerns or values that do not matter in the present moment, and once we have reminded ourselves of what does matter, we can focus on the value or concern at issue.

The main difficulty that remains is to determine when we should seek to enhance our evaluative understanding using our emotion regulation skills. After all, if we can answer this question, then everything else falls into place. If we know that this is a time to enhance our evaluative understanding using our emotion regulation skills, then we can do so. We can engage in appropriate regulation and reflection, and then we can use the understanding that we gain to form a plan. Once we have a plan, we can use our emotion regulation skills to help

us execute this plan. The only question, then, is when we should seek to enhance our understanding.

Intuitively, we should use our emotion regulation skills to enhance our understanding whenever the value of doing so would outweigh the cost of reflection and regulation. But calculating the cost of reflection would itself require deliberation. So what we need, instead, is a set of habits, plans, and policies that would allow us to identify, without deliberation, when reflecting on our present situation is worthwhile. Developing such habits, plans, and policies may initially require practice and instruction, whether from parents, therapists, or other knowledgeable individuals. And, for those who begin this process with poor habits (most likely involving too little reflection), straightening the proverbial stick may require bending it in the opposite direction first, in good Aristotelian spirit. But once in place, these habits can set the correct tone for the individual's emotion-regulatory tendencies.

So what sort of habits will these be? When, in general, should we be disposed to use our emotion regulation skills to enhance our evaluative understanding? Well, in general, regulation and reflection will not be worthwhile in situations requiring urgent action, or in situations where the individual is too tired or unmotivated to reflect, or in situations where spontaneity would serve just as well, because the choice to be made is not significant, according to prior understanding. Reflection will also not be worth it in situations that are already very well understood. But reflection is generally valuable when there is time to reflect, when the individual has sufficient energy and motivation to reflect, when the issue at stake is significant, when it is poorly understood, and when the individual in question is suffering. For in such situations, the value of enhancing one's evaluative understanding and of identifying a calmer,

alternative response is very great, and this value will typically outweigh the cost of reflection and regulation.

The challenge, then, is for individuals is to identify and implement good habits of reflection, and to learn how to properly regulate their emotions, first for the sake of evaluative understanding, and then for the sake of acting in accordance with their understanding. Ideally, individuals will be guided towards these habits and skills by knowledgeable teachers, for achieving them on one's own is no easy task.

But if we can achieve such appropriate habits of reflection and acquire such appropriate regulatory skills, we will be better able to live our lives in a way that is guided by understanding. And this, to borrow Aristotle's phrase, would make "a very great difference, or rather *all* the difference."<sup>2</sup>

---

<sup>2</sup> Aristotle (1991, *NE* 2.1, 1103b27-1104a9).

## Bibliography

- Aristotle. (1984). *Aristotle: The Complete Works*. (J. Barnes, Ed.). Princeton, NJ: Princeton University press.
- Bahlul, R. (2015). Emotion as pathception. *Philosophical Explorations*, 18(1), 104–122. <http://doi.org/10.1080/13869795.2013.874494>
- Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54(7), 462–479. <http://doi.org/10.1037/0003-066X.54.7.462>
- Barlassina, L., & Newen, A. (2014). The Role of Bodily Perception in Emotion: In Defense of an Impure Somatic Theory. *Philosophy and Phenomenological Research*, 89(3), 637–678. <http://doi.org/10.1111/phpr.12041>
- Barrett, L. F., Mesquita, B., Ochsner, K. N., & Gross, J. J. (2007). The Experience of Emotion. *Annual Review of Psychology*, 58(1), 373–403. <http://doi.org/10.1146/annurev.psych.58.110405.085709>
- Baumberger, C., Beisbart, C., & Brun, G. (2017). What is Understanding? An Overview of Recent Debates in Epistemology and Philosophy of Science. In *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. New York, NY: Routledge.
- Beck, A. (1967). *Depression: Clinical, Experimental, and Theoretical Aspects*. Philadelphia, PA: University of Pennsylvania Press.
- Beck, A. (1976). *Cognitive Therapy and the Emotional Disorders*. Oxford, England: International Universities Press.
- Bell, M. (2005). A Woman's Scorn: Toward a Feminist Defense of Contempt as a Moral Emotion. *Hypatia*, 20(4), 80–93.
- Berninger, A. (2016). Thinking sadly: In favor of an adverbial theory of emotions. *Philosophical Psychology*, 29(6), 799–812. <http://doi.org/10.1080/09515089.2016.1159294>
- Boghossian, P. A. (2003). The Normativity of Content. *Philosophical Issues*, 13(1), 31–45.
- Bok, H. (1998). *Freedom and responsibility*. Princeton, NJ: Princeton University press.

- Brady, M. S. (2013). *Emotional Insight*. Oxford, UK: Oxford University Press.
- Carlsson, A. B. (2017). Blameworthiness as Deserved Guilt. *Journal of Ethics*, 21(1), 89–115. <http://doi.org/10.1007/s10892-016-9241-x>
- Carroll, N. (1990). *The Philosophy of Horror, or, Paradoxes of the Heart*. New York, NY: Routledge.
- Carruthers, P. (2017). Valence and Value. *Philosophy and Phenomenological Research*, 1–23. <http://doi.org/10.1111/phpr.12395>
- Chappell, R. Y. (2012). Fittingness: The Sole Normative Primitive. *The Philosophical Quarterly*, 62(249), 684–704.
- Cholbi, M. (2017). Finding the Good in Grief: What Augustine Knew that Meursault Couldn't. *Journal of the American Philosophical Association*, 3, 91–105.
- Cholbi, M. (2017). Grief's rationality, backward and forward. *Philosophy and Phenomenological Research*, XCIV(2), 285–300. <http://doi.org/10.1111/phpr.12353>
- Clarke, R. (2012). Some theses on desert. *Philosophical Explorations*, 16(2), 153–164.
- Cullity, G. (2018). Weighing Reasons. In D. Star (Ed.), *The oxford handbook of reasons and normativity* (pp. 423–442). Oxford, UK: Oxford University Press.
- Cushman, F., Young, L., & Greene, J. D. (2010). Our multi-system moral psychology: Towards a consensus view. In J. Doris, G. Harman, S. Nichols, J. Prinz, W. Sinnott-Armstrong, & S. Stich (Eds.), *The Oxford Handbook of Moral Psychology*. Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199582143.003.0003>
- D'Arms, J., & Jacobson, D. (2000). The moralistic fallacy: On the 'appropriateness' of emotions. *Philosophy and Phenomenological Research*, 61(1), 65–90.
- D'Arms, J., & Jacobson, D. (2003). The significance of recalcitrant emotion (or, anti-quasijudgmentalism). *Royal Institute of Philosophy Supplement*, 52, 127–145. <http://doi.org/10.1017/S1358246100007931>
- Damasio, A. (1994). *Descartes' Error*. New York, NY: Penguin Books.
- Dancy, J. (2004). *Ethics without principles*. Oxford, UK: Oxford University Press.
- De Raedt, R., Baert, S., Demeyer, I., Goeleven, E., Raes, A., Visser, A., ... Speckens, A. (2012). Changes in attentional processing of emotional information following mindfulness-based cognitive therapy in people with a history of depression: Towards an open attention for all

emotional experiences. *Cognitive Therapy and Research*, 36(6), 612–620.  
<http://doi.org/10.1007/s10608-011-9411-x>

de Sousa, R. (1987). *The rationality of emotion*. Cambridge, MA: MIT Press.

Deonna, J. A., & Teroni, F. (2012). From Justified Emotions to Justified Evaluative Judgements. *Dialogue*, 51(01), 55–77. <http://doi.org/10.1017/S0012217312000236>

Deonna, J. A., & Teroni, F. (2015). Emotions as Attitudes. *Dialectica*, 69(3), 293–311.  
<http://doi.org/10.1111/1746-8361.12116>

Deonna, J. A., & Teroni, F. (2016). Value and Emotion. In *Handbook of Value: Perspectives from Economics, Neuroscience, Philosophy, Psychology and Sociology* (pp. 155–174). Oxford, UK: Oxford University Press.

Deonna, J., & Teroni, F. (2013). What role for emotions in well-being? *Philosophical Topics*, 41(1), 123–142.

Dokic, J., & Lemaire, S. (2013). Are emotions perceptions of value? *Canadian Journal of Philosophy*, 43(2), 227–247. <http://doi.org/10.1080/00455091.2013.826057>

Döring, S. A. (2007). Seeing what to do: Affective perception and rational motivation. *Dialectica*, 61(3), 363–394. <http://doi.org/10.1111/j.1746-8361.2007.01105.x>

Elgin, C. Z. (2008). Emotion and Understanding. In G. Brun, U. Doguoglu, & D. Kuenzle (Eds.), *Epistemology and Emotions*. Burlington, VT: Ashgate Publishing.

Ellis, A., & Dryden, W. (1997). *The Practice of Rational Emotive Behavior Therapy*. New York, NY: Springer Publishing.

Franklin, C. E. (2013). Valuing Blame. In D. J. Coates & N. A. Tognazzini (Eds.), *Blame: its nature and norms* (pp. 207–223). New York, NY: Oxford University Press.

Garland, E. L., & Howard, M. O. (2013). Mindfulness-oriented recovery enhancement reduces pain attentional bias in chronic pain patients. *Psychotherapy and Psychosomatics*, 82(5), 311–318. <http://doi.org/10.1159/000348868>

Garland, E. L., Farb, N. A., Goldin, P. R., & Fredrickson, B. L. (2015). The Mindfulness-to-Meaning Theory: Extensions, Applications, and Challenges at the Attention–Appraisal–Emotion Interface. *Psychological Inquiry*, 26(4), 377–387.  
<http://doi.org/10.1080/1047840X.2015.1092493>

Garland, E. L., Farb, N. A., R. Goldin, P., & Fredrickson, B. L. (2015). Mindfulness Broadens Awareness and Builds Eudaimonic Meaning: A Process Model of Mindful Positive Emotion

- Regulation. *Psychological Inquiry*, 26(4), 293–314.  
<http://doi.org/10.1080/1047840X.2015.1064294>
- Garland, E. L., Gaylord, S. A., & Fredrickson, B. L. (2011). Positive Reappraisal Mediates the Stress-Reductive Effects of Mindfulness: An Upward Spiral Process. *Mindfulness*, 2(1), 59–67. <http://doi.org/10.1007/s12671-011-0043-8>
- Garland, E. L., Hanley, A., Farb, N. A., & Froeliger, B. (2015). State Mindfulness During Meditation Predicts Enhanced Cognitive Reappraisal. *Mindfulness*, 6(2), 234–242. <http://doi.org/10.1007/s12671-013-0250-6>
- Garland, E. L., Kiken, L. G., Faurot, K., Palsson, O., & Gaylord, S. A. (2017). Upward Spirals of Mindfulness and Reappraisal: Testing the Mindfulness-to-Meaning Theory with Autoregressive Latent Trajectory Modeling. *Cognitive Therapy and Research*, 41(3), 381–392. <http://doi.org/10.1007/s10608-016-9768-y>
- Garland, E., Gaylord, S., & Park, J. (2009). The Role of Mindfulness in Positive Reappraisal. *Explore: The Journal of Science and Healing*, 5(1), 37–44. <http://doi.org/10.1016/j.explore.2008.10.001>
- Gertken, J., & Kieseewetter, B. (2017). The right and the wrong kind of reasons. *Philosophy Compass*, 12(5), 1–20.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: a theory of normative judgment*. Oxford, UK: Oxford University Press.
- Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues*, 15(1), 338–350.
- Goldie, P. (2008). Misleading Emotions. In G. Brun, U. Doguoglu, & D. Kuenzle (Eds.), *Epistemology and Emotions*. Burlington, VT: Ashgate Publishing.
- Golubickis, M., Tan, L. B. G., Falben, J. K., & Macrae, C. N. (2016). The observing self: Diminishing egocentrism through brief mindfulness meditation. *European Journal of Social Psychology*, 46(4), 521–527. <http://doi.org/10.1002/ejsp.2186>
- Graber ML, Franklin N, & Gordon R. (2005). Diagnostic error in internal medicine. *Archives of Internal Medicine*, 165(13), 1493–1499. <http://doi.org/10.1001/archinte.165.13.1493>
- Greenspan, P. (1988). *Emotions & Reasons*. New York, NY: Routledge.
- Greenspan, P. (2004). Practical Reasoning and Emotion. In A. R. Mele & P. Rawling (Eds.), *The oxford handbook of rationality*. Oxford, UK: Oxford University Press.

- Grimm, S. R. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science*, 57(3), 515–535. <http://doi.org/10.1093/bjps/axl015>
- Gross, J. J. (2014). Emotion regulation: conceptual and empirical foundations. In *Handbook of Emotion Regulation* (pp. 3–22). New York, NY: Guilford Press.
- Grossmann, I., & Kross, E. (2014). Exploring Solomon's Paradox: Self-Distancing Eliminates the Self-Other Asymmetry in Wise Reasoning About Close Relationships in Younger and Older Adults. *Psychological Science*, 25(8), 1571–1580. <http://doi.org/10.1177/0956797614535400>
- Gyurak, A., Gross, J. J., & Etkin, A. (2011). Explicit and implicit emotion regulation: A dual-process framework. *Cognition and Emotion*, 25(3), 400–412. <http://doi.org/10.1080/02699931.2010.544160>
- Hafenbrack, A. C., Kinias, Z., & Barsade, S. G. (2014). Debiasing the Mind Through Meditation: Mindfulness and the Sunk-Cost Bias. *Psychological Science*, 25(2), 369–376. <http://doi.org/10.1177/0956797613503853>
- Haidt, J. (2001). The emotion dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review*, 108(1), 83–95. <http://doi.org/10.1037//0033-295X>.
- Hanley, A. W., & Garland, E. L. (2014). Dispositional mindfulness co-varies with self-reported positive reappraisal. *Personality and Individual Differences*, 66, 46–152. <http://doi.org/10.1016/j.paid.2014.03.014>
- Hanley, A., Garland, E. L., & Black, D. S. (2014). Use of Mindful Reappraisal Coping Among Meditation Practitioners. *Journal of Clinical Psychology*, 70(3), 294–301. <http://doi.org/10.1002/jclp.22023>
- Heuer. (2010). Beyond Wrong Reasons: the buck-passing account of value. In M. Brady (Ed.), *New Waves in Metaethics* (pp. 166–184). Basingstoke: Palgrave-Macmillan.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9), 437–457.
- Hieronymi, P. (2013). The Use of Reasons in Thought (and the Use of Earmarks in Arguments). *Ethics*, 124(1), 114–127. <http://doi.org/10.1086/671402>
- Hookway, C. (2003). Affective States and Epistemic Immediacy. *Metaphilosophy*, 34(1–2), 78. <http://doi.org/10.1111/1467-9973.00261>
- Hopthrow, T., Hooper, N., Mahmood, L., Meier, B. P., & Weger, U. (2017). Mindfulness reduces the correspondence bias. *Quarterly Journal of Experimental Psychology*, 70(3), 351–360. <http://doi.org/10.1080/17470218.2016.1149498>



- Howard, C. (n.d.). The Fundamentality of Fit. In *Oxford Studies in Metaethics*.
- Howard, C. M. (n.d.). Fittingness. *Philosophy Compass*.
- Hume, D. (1739). *A treatise of human nature*. (P. Nidditch & L. Selby-Bigge, Eds.) (1978th ed.). Clarendon Press.
- Hume, D. (1985). Of Tragedy. In *Essays moral, political, and literary*. Indianapolis, IN: Liberty Fund.
- Jaggar, A. M., & Jaggar, A. M. (1989). Love and knowledge: Emotion in feminist epistemology. *Inquiry*, 32(2), 151–176. <http://doi.org/10.1080/00201748908602185>
- Jollimore, T. (2004). Meaningless Happiness and Meaningful Suffering. *Southern Journal of Philosophy*, 42(3), 333.
- Kagan, S. (1992). The limits of well-being. *Social Philosophy and Policy*, 9(2), 169–189.
- Kauppinen, A. (2017). Valuing Anger. In M. Cherry & O. Flanagan (Eds.), *The moral psychology of anger* (pp. 31–48). London: Rowman & Littlefield International Ltd.
- Kiken, L. G., & Shook, N. J. (2011). Looking up: Mindfulness increases positive judgments and reduces negativity bias. *Social Psychological and Personality Science*, 2(4), 425–431. <http://doi.org/10.1177/1948550610396585>
- Kriegel, U. (2014). Towards a new feeling theory of emotion. *European Journal of Philosophy*, 22(3), 420–442. <http://doi.org/10.1111/j.1468-0378.2011.00493.x>
- Krosgaard, C. M. (1983). Two distinctions in goodness. *Philosophical Review*, 92(2), 133-169–195.
- Kross, E., & Ayduk, O. (2008). Facilitating adaptive emotional analysis: Distinguishing distanced-analysis of depressive experiences from immersed-analysis and distraction. *Personality and Social Psychology Bulletin*, 34(7), 924–938. <http://doi.org/10.1177/0146167208315938>
- Kross, E., & Ayduk, O. (2011). Making meaning out of negative experiences by self-distancing. *Current Directions in Psychological Science*, 20(3), 187–191. <http://doi.org/10.1177/0963721411408883>
- Kross, E., & Ayduk, O. (2017). *Self-Distancing: Theory, Research, and Current Directions*. *Advances in Experimental Social Psychology* (1st ed., Vol. 55). Elsevier Inc. <http://doi.org/10.1016/bs.aesp.2016.10.002>

- Kross, E., Ayduk, O., & Mischel, W. (2005). When Asking “Why” Does Not Hurt. *Psychological Science (Wiley-Blackwell)*, 16(9), 709–715. <http://doi.org/10.1111/j.1467-9280.2005.01600.x>
- Kurth, C. (2018). Anxiety: a case study on the value of negative emotion. In *Shadows of the soul: philosophical perspectives on negative emotions* (pp. 95–104). New York, NY: Routledge.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge, UK: Cambridge University Press.
- Lance, M., & Tanesini, A. (2004). Emotion and Rationality. *Canadian Journal of Philosophy*, 34(sup1), 275–295. <http://doi.org/10.1080/00455091.2004.10717608>
- Lang, G. (2008). The right kind of solution to the wrong kind of reason problem. *Utilitas*, 20(4), 472–489. <http://doi.org/10.1017/S0953820809003501>
- Lazarus, R. S. (1991). *Emotion and Adaptation*. New York, NY: Oxford University Press.
- Lueke, A., & Gibson, B. (2015). Mindfulness Meditation Reduces Implicit Age and Race Bias: The Role of Reduced Automaticity of Responding. *Social Psychological and Personality Science*, 6(3), 284–291. <http://doi.org/10.1177/1948550614559651>
- Lueke, A., & Gibson, B. (2016). Brief mindfulness meditation reduces discrimination. *Psychology of Consciousness: Theory, Research, and Practice*, 3(1), 34–44. <http://doi.org/10.1037/cns0000081>
- McFall, L. (1991). What is wrong with bitterness? In C. Card (Ed.), *Feminist ethics* (pp. 146–160). Lawrence: University of Kansas Press.
- Mchugh, C., & Way, J. (2016). Fittingness First. *Ethics*, 126(3), 575–606.
- Moller, D. (2007). Love and Death. *The Journal of Philosophy*, 104(6), 301–316. <http://doi.org/10.1007/s>
- Morton, A. (2010). Epistemic Emotions. In P. Goldie (Ed.), *The Oxford Handbook of Philosophy of Emotion* (pp. 385–400). Oxford, UK: Oxford University Press.
- Nair, S. (2016). How Do Reasons Accrue? In E. Lord & B. Maguire (Eds.), *Weighing Reasons* (pp. 56–73). New York, NY: Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199315192.003.0003>
- Narayan, U. (1988). Working Together Across Difference : Some Considerations on Emotions and Political Practice. *Hypatia*, 3(2), 31–47.

- Neill, J. O. (1992). The varieties of intrinsic value. *The Monist*, 75(2), 119–137.
- Nolen-Hoeksema, S., Wisco, B. E., & Lyubomirsky, S. (2008). Rethinking Rumination. *Perspectives on Psychological Science*, 3(5), 400–424. <http://doi.org/10.1111/j.1745-6924.2008.00088.x>
- Nussbaum, M. (2001). *Upheavals of Thought: The intelligence of emotions*. Cambridge, UK: Cambridge University Press.
- Olson, J. (2004). Buck-Passing and the Wrong Kind of Reasons. *The Philosophical Quarterly*, 54(215), 295–300. <http://doi.org/10.1111/j.0031-8094.2004.00353.x>
- Parfit, D. (2001). Rationality and reasons. In D. Egonsson, J. Josefsson, B. Petersson, & T. Rønnow-Rasmussen (Eds.), *Exploring practical philosophy: From action to values* (pp. 17–39). Ashgate: Aldershot.
- Parfit, D. (2011). *On What Matters*. Oxford, UK: Oxford University Press.
- Pettigrove, G., & Tanaka, K. (2014). Anger and moral judgment. *Australasian Journal of Philosophy*, 92(2), 269–286. <http://doi.org/10.1080/00048402.2013.795990>
- Piller, C. (2006). Content-Related and Attitude-Related Reasons for Preferences. *Royal Institute of Philosophy Supplements*, 81(59), 155–181.
- Prinz, J. (2004). *Gut reactions: a perceptual theory of emotion*. New York, NY: Oxford University Press.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2000). A distinction in value: intrinsic and for its own sake. *Proceedings of the Aristotelian Society*, 100(1), 33–51.
- Rabinowicz, W., & Ronnow-Rasmussen, T. (2004). The Strike of the Demon: On Fitting Pro-attitudes and Value. *Ethics*, 114(3), 391–423. <http://doi.org/10.1086/381694>
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2012). Ethics Discussion at PEA Soup: Rabinowicz and Rønnow-Rasmussen on Way. Retrieved from <http://peasoup.typepad.com/peasoup/2012/06/ethics-discussions-at-pea-soup-rabinowicz-and-ronnow-rasmussen-on-way.html>
- Railton, P. (1994). Truth, Reason, and the Regulation of Belief. *Philosophical Issues*, 5(1994), 71–93.
- Railton, P. (2014). The Affective Dog and Its Rational Tale : Intuition and Attunement. *Ethics*, 124(4), 813–859. <http://doi.org/10.1086/675876>

- Railton, P. (2017). At the Core of Our Capacity to Act for a Reason: The Affective System and Evaluative Model-Based Learning and Control. *Emotion Review*, 9(4), 335–342. <http://doi.org/10.1177/1754073916670021>
- Ranney, R. M., Bruehlman-Senecal, E., & Ayduk, O. (2017). Comparing the Effects of Three Online Cognitive Reappraisal Trainings on Well-Being. *Journal of Happiness Studies*, 18(5), 1319–1338. <http://doi.org/10.1007/s10902-016-9779-0>
- Roberts, R. C. (1988). What an Emotion is: A Sketch. *Philosophical Review*, 97(2), 183–209.
- Roberts, R. C. (2003). *Emotions: an essay in aid of moral psychology*. Cambridge, UK: Cambridge University Press.
- Rorty, A. O. (1980). Explaining Emotions. In A. O. Rorty (Ed.), *Explaining Emotions* (pp. 103–126). Berkeley, CA: University of California Press.
- Rowland, R. (2017). Reasons or Fittingness First ? *Ethics*, 128(1), 212–229.
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110(1), 145–172. <http://doi.org/10.1037/0033-295X.110.1.145>
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge, MA: Harvard University Press.
- Scanlon, T. M. (2013). Giving desert its due. *Philosophical Explorations*, 16(2), 101–116. <http://doi.org/10.1080/13869795.2013.787437>
- Schroeder, M. (2007). *Slaves of the passions*. Oxford, UK: Oxford University Press.
- Schroeder, M. (2010). Value and the right kind of reason. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 5, pp. 25–55). New York, NY: Oxford University Press. Retrieved from [http://www-bcf.usc.edu/~maschroe/research/Schroeder\\_V%26RKR.pdf](http://www-bcf.usc.edu/~maschroe/research/Schroeder_V%26RKR.pdf)
- Schroeder, M. (2012). The Ubiquity of State-Given Reasons. *Ethics*, 122(3), 457–488. <http://doi.org/10.1086/664753>
- Schroeter, L., Schroeter, F., & Jones, K. (2015). Do Emotions Represent Values? *Dialectica*, 69(3), 357–380. <http://doi.org/10.1111/1746-8361.12108>
- Shah, N., & Velleman, J. D. (2005). Philosophical Review Doxastic Deliberation Doxastic Deliberation. *Philosophical Review*, 114(4), 497–534.
- Skorupski, J. (2007). Buck-Passing about Goodness. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, & D. Egonsson (Eds.), *Hommage à Wlodek: Philosophical Papers Dedicated to*

Wlodek Rabinowicz (pp. 1–15). Retrieved from  
<http://www.fil.lu.se/hommageawlodek/site/papper/SkorupskiJohn.pdf>

Skorupski, J. (2010). *The domain of reasons*. Oxford, UK: Oxford University Press.

Smith, A. (1759). *The Theory of Moral Sentiments*. (D. D. Raphael & A. L. Macfie, Eds.) (1976th ed.). Oxford: Oxford University Press.

Smuts, A. (2007). The Paradox of Painful Art. *The Journal of Aesthetic Education*, 41(3), 59–76.  
<http://doi.org/10.1353/jae.2007.0029>

Smuts, A. (2009). Art and Negative Affect. *Philosophy Compass*, 4(1), 39–55.  
<http://doi.org/10.1111/j.1747-9991.2008.00199.x>

Solomon, R. (2007). *True to our feelings*. New York, NY: Oxford University Press.

Srinivasan, A. (2018). The Aptness of Anger. *Journal of Political Philosophy*, 26(2), 123–144.  
<http://doi.org/10.1111/jopp.12130>

Srinivasan, A. (2018). The Aptness of Anger. *Journal of Political Philosophy*, 26(2), 123–144.  
<http://doi.org/10.1111/jopp.12130>

Steglich-Petersen, A. (2006). No norm needed: On the aim of belief. *Philosophical Quarterly*, 56(225), 499–516. <http://doi.org/10.1111/j.1467-9213.2006.455.x>

Steglich-Petersen, A. (2009). Weighing the aim of belief. *Philosophical Studies*, 145(3), 395–405.  
<http://doi.org/10.1007/s11098-008-9239-8>

Tappolet, C. (2005). Ambivalent Emotions and the Perceptual Account of Emotions. *Analysis*, 65:3(287), 229–233. Retrieved from  
<http://ezp.slu.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=phl&AN=PHL2074021&site=ehost-live%5Cnhttp://content.ebscohost.com/ezp.slu.edu/ContentServer.asp?T=P&P=AN&K=17518804&S=R&D=a2h&EbscoContent=dGJyMNHX8kSeprY4xNvgOLCmr0yep65Ss>

Tappolet, C. (2011). Values and Emotions: Neo-Sentimentalism's Prospects. In C. Bagnoli (Ed.), *Morality and the Emotions* (pp. 117–134). New York, NY: Oxford University Press.

Tincher, M. M., Lebois, L. A. M., & Barsalou, L. W. (2016). Mindful Attention Reduces Linguistic Intergroup Bias. *Mindfulness*, 7(2), 349–360. <http://doi.org/10.1007/s12671-015-0450-3>

Troy, A. S., Shallcross, A. J., Davis, T. S., & Mauss, I. B. (2013). History of Mindfulness-Based Cognitive Therapy Is Associated with Increased Cognitive Reappraisal Ability. *Mindfulness*, 4(3), 213–222. <http://doi.org/10.1007/s12671-012-0114-5>

- Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A Person-Centered Approach to Moral Judgment. *Perspectives on Psychological Science*, *10*(1), 72–81.  
<http://doi.org/10.1177/1745691614556679>
- Vahid, H. (2006). Aiming at truth: Doxastic vs. epistemic goals. *Philosophical Studies*, *131*(2), 303–335. <http://doi.org/10.1007/s11098-004-7481-2>
- Way, J. (2012). Transmission and the Wrong Kind of Reason. *Ethics*, *122*(3), 489–515.  
<http://doi.org/10.1086/664749>
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, *16*(16), 267–297.
- Whiting, D. (2010). Should I believe the truth? *Dialectica*, *64*(2), 213–224.  
<http://doi.org/10.1111/j.1746-8361.2009.01204.x>
- Whiting, D. (2012). Are emotions perceptual experiences of value? *Ratio*, *25*(1), 93–107.  
<http://doi.org/10.1111/j.1467-9329.2011.00518.x>
- Whiting, D. (2012). Does belief aim (only) at the truth? *Pacific Philosophical Quarterly*, *93*(2), 279–300. <http://doi.org/10.1111/j.1468-0114.2012.01421.x>
- Williams, B. (1973). Deciding to believe. In B. Williams (Ed.), *Problems of the self* (pp. 136–151). Cambridge: Cambridge University Press.
- Yiend, J. (2010). The effects of emotion on attention: A review of attentional processing of emotional information. *Cognition and Emotion*, *24*(1), 3–47.  
<http://doi.org/10.1080/02699930903205698>
- Zagzebski, L. (2001). Recovering understanding. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. New York, NY: Oxford University Press.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, *35*(2), 151–175. <http://doi.org/10.1037/0003-066X.35.2.151>
- Zajonc, R. B. (1985). The primacy of affect. *American Psychologist*, *40*(7), 849–850.  
<http://doi.org/10.1037/0003-066X.40.7.849>