



Harrison, Eilidh (2021) *Affective justification: how emotional experience can epistemically justify evaluative belief*. PhD thesis.

<http://theses.gla.ac.uk/82279/>

-

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>  
[research-enlighten@glasgow.ac.uk](mailto:research-enlighten@glasgow.ac.uk)

**Affective Justification: How Emotional Experience Can Epistemically Justify  
Evaluative Belief**

Eilidh Harrison  
MA (Hons), MSc

Submitted in fulfilment of the requirements for the Degree of  
Doctor of Philosophy

Philosophy  
School of Humanities  
College of Arts  
University of Glasgow

March 2021

## Abstract

The idea that emotional experience is capable of lending immediate prima facie epistemic justification to evaluative belief has been amassing significant philosophical support in recent years. The proposal that it is my anger, say, that justifies my belief that I've been wronged putatively provides us with an intuitive and naturalised explanation as to how we receive immediate and defeasible justification for our evaluative beliefs. With many notable advocates in the literature, this justificatory thesis of emotion is fast becoming a central facet in how we conceive of the emotions' epistemic role with respect to our everyday lives.

Interestingly, however, despite the fact that the justificatory thesis is fundamentally an epistemological proposal, comparatively little of the philosophical literature has been dedicated to exploring the epistemological avenues through which emotions might be capable of delivering such an epistemic yield. Accordingly, the central purpose of this thesis is to provide a novel and thorough analysis of *how* emotional experience might be capable of playing this justificatory role. Here, I present and evaluate three broad models of emotional justification: emotional dogmatism, emotional reliabilism, and agent-based views. Emotional dogmatist views, I argue, fail in virtue of being vulnerable to over-generalisation worries and problematic commitments to the contents of emotional awareness. Emotional reliabilism, while possessing the resources to avoid some objections, is vulnerable to worrisome clairvoyance-style challenges which establish the insufficiency of emotional reliability for epistemic justification. Finally, having learned our lessons from the shortcomings of these views, I argue that an agent-based theory grounded in the development of learned emotional competences provides the most plausible account of how emotional experience can epistemically justify evaluative belief. This discussion, I believe, will both illuminate contemporary discussions of the justificatory thesis of emotion found in the literature, and provide novel insight into the epistemic capacities of the emotions.

## Table of Contents

<b>Acknowledgements.....</b>	<b>5</b>
<b>Author’s Declaration.....</b>	<b>6</b>
<b>Chapter 1: Introducing the Justificatory Thesis of Emotion .....</b>	<b>7</b>
1.1 Introduction.....	7
1.2 Preliminaries: Emotional Experience.....	11
1.3 Preliminaries: Epistemic Justification.....	13
1.4 Preliminaries: Evaluative Belief.....	15
1.5 The Structure of the Dissertation.....	16
<b>Chapter 2: A Dogmatist Approach.....</b>	<b>20</b>
2.1 Introductory Remarks.....	20
2.2 Basic Emotional Dogmatism.....	22
2.2.1 Objection: An Over-Generalisation Problem.....	24
2.3 Restricted Emotional Dogmatism.....	32
2.4 Objection: Evaluative Properties and Emotional Awareness.....	36
2.4.1 Evaluative Properties as Direct Objects of Awareness.....	41
2.4.2 Evaluative Properties as Indirect Objects of Awareness.....	53
2.5 Concluding Remarks.....	55
<b>Chapter 3: The Problem with Restricted Emotional Dogmatism.....</b>	<b>58</b>
3.1 Introductory Remarks.....	58
3.2 Objection: The Dilemma of Evaluative Truth-Makers.....	59
3.2.1. Truth-Makers as Evaluative Properties.....	61
3.2.2. Truth-Makers as Non-Evaluative Properties.....	67
3.3 Alternative Restricted Views.....	71
3.3.1. Receptive Seemings Emotional Dogmatism.....	72
3.3.2. Knowledge-How Emotional Dogmatism.....	76
3.4 Summing Up.....	80
<b>Chapter 4: A Reliabilist Approach.....</b>	<b>82</b>
4.1 Introductory Remarks.....	82
4.2 Motivating and Clarifying Emotional Reliabilism.....	83
4.2.1 A Causal Problem.....	86
4.2.2 A Generality Problem.....	88

4.2.3 A Necessity Problem.....	91
4.3 Objection: The Unreliability Problem.....	95
4.3.1 The Cognitive Penetrability of Emotion.....	96
4.3.2 The Obscurity of Value.....	101
4.3.3 Against Emotional Reliability in Virtuous Agents.....	103
4.4 Concluding Remarks.....	109
<b>Chapter 5: The Problem with Emotional Reliabilism.....</b>	<b>112</b>
5.1 Introductory Remarks.....	112
5.2 Objection: The Insufficiency of Emotional Reliability .....	113
5.3 Modifying Emotional Reliabilism.....	122
5.3.1 Evidentialist Emotional Reliabilism.....	122
5.3.2 Agent Emotional Reliabilism.....	127
5.4 Summing Up.....	130
<b>Chapter 6: Virtue, Character, and Competence: Agent-Based Approaches.....</b>	<b>132</b>
6.1 Taking Stock .....	132
6.2 Virtue Epistemology and Emotional Justification .....	135
6.2.1 Virtue Responsibilism.....	136
6.2.2 Virtue Reliabilism.....	140
6.3 The Emotional Competence View.....	144
6.4 Objection: An Ambiguity in the Scope of Emotional Competences .....	147
6.5 Solution: Emotional Competences Localised to Eliciting Scenarios.....	152
6.5.1 The Processes of Emotional Learning.....	155
6.5.2 Emotional Learning and the Insight from Oppression.....	161
6.6 The Virtues of the Learned Emotional Competence View.....	169
6.7 Concluding Remarks.....	173
<b>Conclusion.....</b>	<b>175</b>
<b>Bibliography.....</b>	<b>179</b>

## **Acknowledgements**

I owe a great deal of gratitude to my supervisors, Michael Brady and Robert Cowan, for their guidance and encouragement throughout my time at Glasgow. It's difficult to verbalise just how much my academic life has been shaped by their influence. I'm thankful for all they've done to help bring me here. Thanks also to the broader philosophical community at Glasgow for their willingness to give help and advice whenever it was needed.

Beyond Glasgow, my work has been significantly improved by the constructive criticism that I've received from a vast number of conference delegates and from anonymous referees at Synthese and Philosophical Studies. I also gratefully acknowledge the financial support that I received for this research from the College of Arts (2017 – 2020) and the Royal Institute of Philosophy (2020 – 2021).

I'm hugely grateful to Adriana Alcaraz-Sánchez, Laura Fearnley, Nathan Kirkwood, and María Pía Méndez Mateluna for their warmth and friendship over the years, and for enabling my compulsion to spend all day in local coffee shops instead of working in the office. Thanks especially to Nathan for weathering it all with me. Finally, thanks to my wonderful parents, Gillian and Keith, and to the rest of my family for their limitless love and support, and for patiently listening to me complain about this thesis throughout the last four years. The greatest gift I can give them all in return, I'm sure, is never mentioning it ever again.

## **Author's Declaration**

I confirm that this thesis is my own work and that I have: (i) read and understood the University of Glasgow Statement on Plagiarism, (ii) clearly referenced, in both text and the bibliography or references, all sources used in the work; (iii) fully referenced (including page numbers) and used inverted commas for all text quoted from books, journals, web, etc.; (iv) provided the sources for all tables, figures, data, etc. that are not my own work; (v) not made use of the works of any other student(s) past or present without acknowledgement. This includes any of my own works, that has been previously, or concurrently, submitted for assessment, either at this or any other educational institution; (vi) not sought or used the services of any professional agencies to produce this work; (vii) in addition, I understand that any false claim in respect of this work will result in disciplinary action in accordance with University regulations.

I declare I am aware of and understand the University's policy on plagiarism and I certify that this thesis is my own work, except where indicated by referencing, and that I followed the good academic practices noted above.

## Chapter 1

### Introducing the Justificatory Thesis of Emotion

#### 1.1 Introduction

Emotional experience is a significant constituent of our mental lives. Our emotional reactions and experiences perform a variety of valuable roles throughout the course of our lifetimes. Emotions mark and characterise our relationships with others, instigate and sway our processes of decision-making, prompt moral assessments and judgments of ourselves and those around us, facilitate our aesthetic engagement with the arts, and so forth. Given the power, salience, and ubiquity of emotional influence, then, it is unsurprising that the experience of emotion has attracted significant philosophical attention in recent years.

One particularly interesting dimension of the emotions' significance is their *epistemic* import. Emotional experience seems to illuminate the evaluative world around us. We frequently form evaluative beliefs about surrounding objects and states of affairs on the basis of our emotional experiences. I might believe that, say, an offhand remark made by my colleague is offensive in light of my anger in response to it, or that the view from the summit of Ben Lomond is beautiful on the basis of my experience of awe before it, or that a particular joke is amusing on the basis of my experience of amusement, and so forth. Our beliefs about values, like the admirable, the disgusting, the contemptible, the amusing, the enviable, etc. thereby appear to be intimately tied up with and influenced by our emotional experiences.

Now, assuming that one of our epistemic goals is to form justified beliefs about value, an immediate question that arises here is whether the emotions' influence over the formation of our evaluative beliefs is, on balance, for better or for worse. Does perceiving value through the emotional lens tend to help or hinder our pursuit of justified evaluative belief? Does the influence of anger, say, typically aid in or frustrate the formation of justified evaluative beliefs about what is offensive? In recent years, the tide has been turning in favour of positive answers to these questions. Increasingly, many have endorsed the idea that emotional experience bears positive epistemic status insofar as it provides us with epistemic



justification for our evaluative beliefs. Accordingly, the specific view that will be evaluated in this dissertation is what will be referred to as *the justificatory thesis of emotion*.<sup>1</sup>

The Justificatory Thesis of Emotion (JTE): emotional experience is capable of lending immediate *prima facie* epistemic justification to evaluative belief.

Put roughly, what JTE amounts to is the claim that emotional experiences have the ability to act as defeasible reasons in favour of an agent's holding the relevant evaluative belief. On this view, my emotional experience of anger towards my colleague's offhand remark is capable of defeasibly justifying my evaluative belief that the remark was offensive; my guilt in response to my breaking a promise to a friend is capable of defeasibly justifying my evaluative belief that my actions were morally wrong.

Many philosophers have endorsed JTE as a significant component of their account of the epistemological value of emotional experience.<sup>2</sup> Tappolet (2016), for instance, states that "it appears plausible to claim that when you feel the emotion of fear, say, this not only prompts you to believe that what you are afraid of is fearsome, but you are also *prima facie* justified in believing that what you are afraid of is fearsome" (p. 168). Likewise, Döring (2003) argues that "an emotion... resembles a sense-perception in having an intentional content that is representational... as a consequence, an emotion can justify a belief" (p. 215), while Elgin (2008) suggests something close to JTE insofar as she argues that emotional experiences are "initially tenable" in virtue of the fact that "an agent, in the grip of an emotion, has a tendency to credit its deliverances" (p. 34).

There are a number of reasons as to why JTE, if true, is an important claim. The first of which is that JTE stands in opposition to a sceptical view concerning the epistemic capacities of the emotions, namely, the view that emotional experiences cloud our evaluative perception and make us epistemically worse-off than we would be if we were free from emotional influence.<sup>3</sup> If it is the case that emotional experience instead constitutes a defeasible reason for holding particular evaluative beliefs, then JTE supports the conclusion that emotions are

---

<sup>1</sup> I'm borrowing this label from Pelser (2011) and (2014).

<sup>2</sup> Although, see Brady (2013) for notable criticism of the thesis.

<sup>3</sup> Notable philosophical instances of similar sceptical views include the Stoics' assertion that emotions (suitably qualified) "imply false judgments" (Graver 2007, p. 5) or the Kantian idea that affective experiences are "sudden and rash, making reflection impossible" (Borges 2004, p. 144).

best considered as being part and parcel of the rational agent's cognitive architecture, rather than psychological pollutants that impede our epistemic inquiries. This is an interesting claim, given the traction the sceptical view seems to have in folk discourse about the emotions.

Second, JTE has a certain explanatory power. Recall that Elgin takes emotions to be experienced as "initially tenable" by the subject undergoing the experience. According to Elgin, we're ordinarily inclined to give credence to the beliefs formed on the basis of emotional experiences.<sup>4</sup> If JTE is true, this makes sense of the trust we place in our emotions. Third, JTE can provide us with a non-mysterious epistemology of value. In illustration of this point, contrast JTE with rational intuitionist views, according to which we come to have justified evaluative beliefs through the faculty of evaluative intuition.<sup>5</sup> In accordance with such a view, one might argue that my evaluative belief that deliberately causing harm is morally wrong is justified simply in virtue of my grasping the self-evidence of that proposition. Now, one common objection facing these accounts is that positing rational intuition as the root of our evaluative epistemology objectionably hinges our possession of justified evaluative belief on a mysterious and unsubstantiated faculty.<sup>6</sup> By contrast, JTE need not posit *any* kind of strange evaluative faculty in order to explain how our evaluative beliefs come to be justified; the explanatory question is answered simply and naturally by appeal to emotional experience. Indeed, bestowing emotional experience with justificatory power not only provides us with an intuitive and naturalised explanation as to how our evaluative beliefs can enjoy positive justificatory status, but, given the putative importance of justification for the acquisition of further epistemic goods, JTE may also be able to deliver a substantive epistemic yield which extends beyond justified belief and plausibly into the domain of evaluative knowledge and understanding.<sup>7</sup>

Finally, and on a related note, the focus of JTE on the *immediacy* of the justification conferred by emotional experience is significant given that it putatively provides us with a way out of pernicious sceptical worries pertaining to the acquisition of justified evaluative belief. If, instead, emotional experience could only transfer justification to evaluative belief

---

<sup>4</sup> Döring (2007) and Pelser (2011) also make this point.

<sup>5</sup> See, for example, Ross (1930) and Audi (1998).

<sup>6</sup> One particularly notable version of this argument is Mackie's (1977) 'Argument From Queerness'.

<sup>7</sup> This point assumes that a belief's being epistemically justified is in some way necessary (but not sufficient) for that belief to count as an instance of knowledge or understanding.

*mediated* by the possession of the subject's other justified evaluative beliefs, then it seems that we would open the door to sceptical concerns about how *those* evaluative beliefs end up being justified. Allowing emotional experience to serve as an experiential justificatory bedrock for evaluative belief, so to speak, allows us to sidestep these sceptical worries (more on this immediacy point in §1.3).

So, thus far, we can recognise JTE as a popular view within the philosophy of emotion which enjoys a number of theoretical advantages. However, if the thesis is to be plausible, there remains an essential question which must be answered, namely, the question of *how* emotional experience is capable of justifying evaluative belief. Interestingly, despite the fact that JTE is fundamentally an epistemological proposal, comparatively little of the philosophical literature has been dedicated to exploring the epistemological avenues through which emotions might be capable of delivering such an epistemic yield. That is, little has been done to explicitly isolate and explore *what it is* about emotional experience that bestows it with the ability to justify evaluative belief. I aim to close this gap in the literature by offering a novel and thorough analysis of various accounts of emotional justification, i.e. accounts as to how and when emotional experience can immediately and defeasibly justify evaluative belief.<sup>8</sup>

Put specifically, the purpose of this dissertation is to present and evaluate three broad epistemological models of emotional justification: dogmatist approaches, reliabilist approaches, and agent-based (or virtue epistemological) approaches. Roughly speaking, emotional dogmatism is the view that emotional experience can immediately and defeasibly justify evaluative belief in virtue of its possession of an epistemically significant *phenomenal character*. Emotional reliabilism, on the other hand, states that the emotions' justificatory ability is rooted in the *reliability* of forming evaluative beliefs on the basis of emotional experience. Finally, agent-based views identify features of the emoter herself, such as her intellectual character traits or reliable cognitive faculties, as those things which bestow her emotional experiences with justificatory power. In what follows, I will argue that dogmatist and reliabilist approaches to emotional justification are not capable of providing us with a

---

<sup>8</sup> My use of 'emotional justification' to refer to the justificatory power of emotion with respect to evaluative belief is not to be confused with the discussion of emotions *themselves* being justified (see Echeverri (2019) for an in-depth discussion of emotional justification in this latter sense). 'Emotional justification' will always be used to refer to the former in this thesis.

satisfactory explanation as to how emotional experience can justify evaluative belief. Instead, I will argue in favour of a particular agent-based view which places epistemic significance on the development of learned emotional competences.

In pursuit of these goals, there are a number of preliminary issues that must be addressed. The question of whether and how emotional experience is capable of playing the justificatory role specified above will naturally be contingent on a number of further questions, such as what we take the structure and nature of emotional experience, epistemic justification, and evaluative belief to be. In what follows, I will lay the groundwork for the assumptions that I will be making and the debates that will be bracketed in order to pave the way for the forthcoming discussion. I will then conclude the chapter by further clarifying the structure of the dissertation and presenting a detailed chapter-by-chapter plan.

## 1.2 Preliminaries: Emotional Experience

Before we begin our epistemological investigation into the justificatory prospects of emotional experience, there's one particularly obvious question that must first be addressed. Namely, what *is* an emotional experience?

First, let me be clear that by 'emotional experience', I intend to capture the phenomenon of occurrent emotional episodes rather than longstanding emotional dispositions, sentiments, or moods. Philosophers of emotion tend to agree that these occurrent emotional experiences involve the following features.<sup>9</sup> First, paradigmatic emotional experiences are felt experiences. That is, they have a phenomenal character (i.e. a something-that-it-is-like to undergo the experience), and tend to bring about a range of bodily feelings and physiological changes.<sup>10</sup> Second, emotional experiences are typically experienced passively. We often describe ourselves as being *overcome* by anger, or grief, or joy; it doesn't seem to be the case that we will our emotional experiences into existence. Third, as Döring notes above, emotions are *intentional* mental states. That is, emotions are not experienced as directionless episodes, but are rather are *about* specific objects and events. To be more specific, emotional

---

<sup>9</sup> See Roberts (1988) for a helpful elucidation of these features.

<sup>10</sup> Although, this varies across different emotion-types. The intellectual pleasure of solving a puzzle, for instance, may not involve physiological feelings in the same way or to the same extent that, say, an episode of anger or fear typically does.

experience *represents* the world as being a certain evaluative way, and can be assessed in terms of accuracy conditions accordingly. Fourth, and finally, emotions are often experienced alongside specific evaluative judgments and action tendencies. The anger that I experience in response to my colleague's offhand comment, for example, will typically come alongside the evaluative judgement that the contents of the comment were harmful or incendiary, and the behavioural tendency to confront or reprimand the colleague.<sup>11</sup>

The primary point of contention amongst philosophers of emotion is identifying precisely what an emotional experience *is*. In other words, there's a crucial question of which of these features, if any, the emotional experience is identical to and which of these features are merely concurrent with the experience. Feeling theorists, for example, identify the emotional experience with the felt component. More specifically, feeling theorists argue that the emotion is identical to physiological feelings or perceptions of bodily change.<sup>12</sup> On such a view, my emotion of anger is identical to the feelings of bodily change that I undergo after having heard my colleague's comment, e.g. a quickening heartbeat, a flush of heat, tensed muscles, various visceral sensations, etc. Judgementalists, on the other hand, prioritise the cognitive features of emotional experience and identify the emotion with evaluative judgment(s),<sup>13</sup> such that my anger is to be understood as identical to a particular kind of evaluative judgment that the comment was harmful or offensive. Perceptual theorists, by contrast, advocate on behalf of the view that emotional experiences are best understood as perceptual experiences of value.<sup>14</sup> On this account, my experience of anger is best understood as a perceptual experience of offense, my grief as a perceptual experience of loss, and so forth.

---

<sup>11</sup> Of course, this is not always the case. We can, and seemingly often do, experience emotions which appear to conflict with our considered evaluative judgments (e.g. a recalcitrant fear of a house spider that I judge to be harmless). Likewise, emotions can be experienced in the absence of any obvious behavioural inclinations or tendencies (e.g. perhaps an experience of surprise at a particularly noteworthy piece of news, or an experience of pride after finishing a particularly difficult task).

<sup>12</sup> See James (1884) for one of the most notable presentations of such a view, and Prinz (2004) for a modern defence of a revised feeling theory.

<sup>13</sup> See Solomon (1988) and Nussbaum (2001) for nuanced defences of the judgementalist view.

<sup>14</sup> See Tappolet (2016) and Döring (2007). Perceptual theories are typically motivated by drawing certain parallels between perceptual and emotional experiences, i.e. perceptual and emotional experiences both seem to be intentional and passively experienced mental states which instantiate phenomenological properties.

In this dissertation, I am not concerned with taking a stance on the substantive question of the nature of emotional experience. In the interest of keeping my discussion open to a variety of views on the nature and structure of emotional experience, I will not assume or defend any particular account which attempts to answer the metaphysical question of what an emotional experience *is*. For the purposes of the dissertation, however, note that a key assumption that I will be making is that emotional experience has intentional content (i.e. represents evaluative properties) and phenomenological character. I take these to be fairly uncontroversial assumptions that are made by many in the literature.<sup>15</sup>

### 1.3 Preliminaries: Epistemic Justification

JTE proposes that emotional experience acts as a source of *immediate* and *prima facie* epistemic justification for evaluative belief. In this section, I'll provide a brief elucidation of these epistemological concepts.

First and foremost, the justification that we're interested in here is specifically epistemic.<sup>16</sup> Epistemic justification is a property that is possessed by certain beliefs. Very loosely speaking, attributing the property of epistemic justification to a belief is similar to evaluating that belief as being supported by good reason(s) or evidence. As a further qualification, note that the kind of epistemic justification primarily at issue here is propositional as opposed to doxastic justification. That is, we're interested in the question as to how and when emotional experience can constitute support for believing a particular evaluative proposition, rather than the question of whether a given subject appropriately believes an evaluative proposition on the basis of their emotional experience.

Second, it is important to distinguish between mediate and immediate justification. To illustrate, suppose that I have a visual experience of a small insect perched on my shoulder.

---

<sup>15</sup> Note, however, that Deonna and Teroni (2012a) endorse an attitudinal view of emotional experience which is incompatible with the claim that emotions have evaluative properties as their intentional content.

<sup>16</sup> Epistemic justification can come apart from other types of justification, such as moral or prudential justification. I might be morally justified in believing that, say, my friend won't break her promise to me (assuming that there are strong moral reasons to have faith in one's friends), but not epistemically justified if I lack evidence in favour of her trustworthiness. Alternatively, I might be prudentially justified in believing that, say, my illness will subside (assuming that there are strong prudential reasons to engage in positive thinking), but not epistemically justified if I lack evidence for this prognosis.

As a result of this visual experience, I then come to form the corresponding belief that there is a small insect on my shoulder. Generally speaking, visual experiences are taken to confer *immediate* justification to belief insofar as that justification exists independently of any of the subject's other justified beliefs. In other words, the justification conferred by my visual experience of the insect does not depend on my having the justified belief that, say, my visual system is reliable at tracking small insects, or that I'm not in a sceptical brain-in-the-vat scenario, and so forth. On the other hand, compare this with a scenario in which my friend tells me that there's a small insect on my back. In this case, the insect is beyond my perceptual field such that I can't see or feel it for myself. Thus, my belief that there's a small insect on my back is plausibly only *mediately* justified given that my justification for this belief *does* depend on my having other justified beliefs, e.g. that my friend is a reliable testifier, that she has no reason to lie to me about the presence of small insects, that her visual faculties tend to function correctly, and so forth. What's important for our purposes is JTE's claim that emotional experience is capable of conferring *immediate* justification to evaluative belief. That is, similar to the visual case mentioned above, my experience of anger can justify my evaluative belief that my colleague's comment was offensive *independently* of my possessing any other justified beliefs concerning, for example, the reliability of my emotional and auditory systems.<sup>17</sup>

Third, and finally, note that the type of justification putatively conferred by emotional experience is qualified as *prima facie* (or defeasible) rather than *ultima facie* justification. Reconsider the visual example above. My visual experience of the insect on my shoulder plausibly confers *prima facie* justification to my belief that there is an insect on my shoulder in light of the fallibility of visual systems. Our visual experiences can, and sometimes do, get things wrong. This justification conferred by my visual experience may be quashed by defeating evidence, such as my having recently ingested a hallucinogenic, my being in bad lighting, and so forth. If, however, there is *no* defeating evidence, then my belief about the small insect would plausibly be *ultima facie* justified (i.e. justified all-things-considered) on

---

<sup>17</sup> One objection that may be raised here is that emotional experiences are precluded from conferring immediate justification in light of their dependence on *cognitive bases*, i.e. prior mental states like perceptions, beliefs, memories, imaginative states, and so forth. Given limitations of space, I won't discuss this objection here (although I address a similar point in §3.3.1), but see Cowan (2018) and Milona and Naar (2020) for detailed discussions of this challenge. It's important to note that neither take this challenge to be capable of barring emotional experience from conferring generative or immediate justification to evaluative belief altogether.

the basis of my visual experience. JTE is only concerned with the possibility of emotional experiences conferring prima facie justification to evaluative belief in light of the seemingly analogous fallibility of our emotional systems.

#### **1.4 Preliminaries: Evaluative Belief**

So, thus far, we have two components of JTE spelled out. Emotional experience, as I'll understand it, is occurrent, intentional, and bears phenomenal character. Moreover, the kind of epistemic justification that we're interested in is propositional, immediate, and defeasible. Let me finish these preliminary considerations with some comments on how evaluative belief will be understood throughout this thesis.

Throughout most of the forthcoming discussion, I will not be making any assumptions about the structure of evaluative belief other than the straightforward conception of an evaluative belief consisting in a propositional attitude attributing a particular evaluative property (or properties) towards a particular object or state of affairs. My evaluative belief that my colleague's comment was offensive, then, represents the comment as instantiating the property of 'offensiveness'.

Now, naturally, a question might arise here concerning the metaphysics of value concepts and properties. Recall from our initial discussion in §1.1. that there seems to be a close connection between our evaluative beliefs and our emotional responses. Drawing from this connection, one plausible view concerning the nature of evaluative concepts and properties is that they're essentially *response-dependent*, i.e. they cannot be understood without reference to the responses of observers.<sup>18</sup> This response-dependence claim can be spelled out in a variety of different ways. One might think that evaluative properties merely causally elicit particular emotional responses, such that a state of affairs is shameful if and only if that state of affairs elicits responses of shame in observers under certain conditions. If such a simple subjectivism strikes the reader as implausible, however, it might be thought that evaluative

---

<sup>18</sup> One interesting thing to note is that endorsing a response-dependent construal of evaluative concepts need not preclude one from being a realist about the nature of certain evaluative properties picked out by those concepts. In other words, it's possible to advocate on behalf of a view where the evaluative concept 'admirable' cannot be explained without reference to the admiration responses of observers, but the evaluative property 'admirable' is itself fully objective and non-relational. See, for example, Tappolet's (2016, p. 116) portrayal and defence of Sentimental Realism.



properties do not merely causally elicit certain emotional responses, but instead *dispose* ordinary observers to experience certain emotional responses under certain conditions. Alternatively, others opt for the view that evaluative properties do not just cause or dispose but *merit* certain emotional responses.<sup>19</sup> Call this latter view neo-sentimentalism.<sup>20</sup>

I am not going into this discussion with any pre-existing view in mind as to which of these is the correct view of the nature of evaluative concepts and properties. While Chapter 2 (§2.4) will comment on the compatibility of neo-sentimentalism with a particular version of JTE (namely, a dogmatist view which states that emotional experience brings epistemically significant phenomenological awareness of evaluative properties), I do not intend for the plausibility of the arguments made in this thesis to substantially hinge on any particular preconception of the metaphysics or ontology of value.

## 1.5 The Structure of the Dissertation

Now that we've addressed the relevant preliminary questions, let me spell out the structure of the forthcoming discussion. To reiterate, the aim of this thesis is to clarify and evaluate three broad views of emotional justification. In Chapter 2, I begin by presenting the first view: *emotional dogmatism* (ED). According to ED, S's emotional experience is capable of immediately and defeasibly justifying their evaluative belief that *e* (where *e* signifies a proposition which attributes an evaluative property to a particular object or event) insofar as it makes it seem to S that *e*. This view is based on a phenomenal conservative (PC) account of experiential justification, according to which an experience is capable of justifying one's belief that *p* insofar as it makes it seem to one that *p*. I then argue that ED fails in virtue of lacking the theoretical resources to exclude problematic cases from enjoying its epistemic yield.

I then suggest that a *restricted* emotional dogmatist view (RED), based on Chudnoff's restricted phenomenal conservative account, is better placed as a model for our investigation into the prospects of an emotional dogmatist view. According to this view, S's emotional experience is capable of immediately and defeasibly justifying their evaluative belief *e*

---

<sup>19</sup> Note that Chapter 2 (§2.4) will summarise a very brief argument against mere causal and dispositional theories of value and in favour of neo-sentimentalist views.

<sup>20</sup> For notable discussion of neo-sentimentalism, see D'Arms and Jacobsen (2000).

insofar as it both (i) makes it seem to S that *e* and (ii) makes it seem as if S is aware of a truth-maker for *e*. Having identified RED as the target view, I then defend this account from an objection particularly salient in the literature, namely, that evaluative properties are not suitable objects of emotional awareness. I argue that RED has a (prima facie) plausible story to tell with respect to evaluative properties being both direct *and* indirect objects of emotional awareness, such that the account has the resources to provide a response to this worry.

Having built a strong philosophical foundation for the account, Chapter 3 then proposes what I take to be the most serious objection facing the view. Namely, I argue that RED's inclusion of the truth-maker condition raises a troubling dilemma for the account. Either RED identifies evaluative properties themselves as truth-makers for evaluative propositions, in which case the view will continue to over-generalise, or it identifies the conjunction of non-evaluative properties which gives rise to the relevant evaluative property as the truth-maker, in which case the view ends up being objectionably restrictive. Either way, RED is generating the wrong result. I then conclude the chapter, and thus our investigation into the prospects of a dogmatist approach to emotional justification, by rejecting two alternative instantiations of the restricted emotional dogmatist thesis: Receptive Seemings Emotional Dogmatism (RSED) and Knowledge-How Emotional Dogmatism (KHED).

In Chapter 4, I then present the second contender for a plausible approach to explaining emotional justification: *emotional reliabilism* (ER). According to ER, S's emotional experience is capable of immediately and defeasibly justifying their evaluative belief *e* insofar as S's forming *e* on the basis of that emotional experience is a reliable process of belief formation. The motivation for this view is twofold. First, it carries compelling weight in the literature, and, second, it putatively avoids the kinds of problems faced by the dogmatist approach. Having motivated and clarified the reliabilist view, I then defend it against several objections. Most notably, I defend ER against what I take to be the most conspicuous objection facing it in the literature, namely, that emotional experience is significantly *unreliable* such that it cannot constitute a reliable belief-forming process with respect to an agent's evaluative beliefs. By considering and rejecting what I take to be the three most plausible arguments in favour of this unreliability, I argue that there is no plausible theoretical support for this claim, such that ER can survive the objection unscathed.

In Chapter 5, again, having built ER into the most philosophically robust version it can be, I level a further objection against the view which I take to highlight a significant problem for the reliabilist approach to emotional justification, namely, that reliably forming true evaluative beliefs on the basis of one's emotional experiences is insufficient for the justification of those beliefs. This objection takes the form of an emotional analogue to traditional clairvoyance cases levelled against reliabilist epistemologies. Having done so, I then rebut two possible modifications of the emotional reliabilist thesis, Evidentialist Emotional Reliabilism (EER) and Agent Emotional Reliabilism (AER), and conclude our investigation into the prospects of a reliabilist approach to emotional justification.

In Chapter 6, I begin by taking stock of my discussion thus far and constructing various desiderata for a plausible theory of emotional justification that we can glean from the weaknesses of the dogmatist and reliabilist approaches. From these desiderata, I present the third class of views to be considered in this dissertation, namely, *agent-based* views of emotional justification, i.e. virtue epistemological views which hinge the justificatory power of emotional experience on features of the emoter herself. Within this class, I identify two broad species of agent-based views: virtue responsibilist views and virtue reliabilist views. I argue that, in light of the shortcomings of virtue responsibilism, a competence-based virtue reliabilism inspired by the work of Sosa is best placed for an investigation into the prospects of an analogous view of emotional justification. On this note, I present and motivate Carter's Emotional Competence (EC) view, according to which S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if (i) that experience is formed via the exercise of her generative emotional competence, and (ii) *e* is formed via the exercise of her doxastic emotional competence.

In spite of the view's attractions, I then argue that EC requires further development in virtue of a generality-style ambiguity in the scope of emotional competences. Put roughly, I argue that EC is silent as to whether: (1) one possesses emotional competence *overall*, (2) one possesses emotional competences localised to particular emotion-types, or (3) one possesses emotional competences localised to particular emotion-types which are experienced in response to particular eliciting scenarios. I argue that clarifying this ambiguity is essential for evaluating the plausibility of the view given that, without a clear understanding of the scope of emotional competence, EC will often be unable to provide instructive epistemic verdicts.

In light of this concern, I propose that emotional competence ought to be localised to particular emotion-types experienced in response to particular eliciting scenarios. My reason for this is that, in my view, the most plausible way of explaining the development of emotional competence is by reference to processes of emotional learning which are localised to particular emotion-types and eliciting scenarios. Put roughly, I propose that a subject, through repeated experience with and exposure to a particular kind of eliciting scenario, comes to develop a localised emotional competence with respect to that eliciting scenario by emotionally learning that this scenario instantiates a particular evaluative property. I then suggest that this analysis prompts a revised construal of EC, according to which a subject S's emotional experience which attributes evaluative property E to object O is capable of immediately and defeasibly justifying S's evaluative belief that O is E if and only if (i) S has the emotion as a result of her having emotionally learned that O is E, and (ii) S forms the belief that O is E on the basis of that emotion. I then defend this view by further clarifying the processes and significance of emotional learning, highlighting a unique selling point of the view with respect to a compelling insight from feminist philosophy, and summarising how the view meets the desiderata previously identified for a plausible theory of emotional justification. Finally, I close the thesis with a brief review of the critical points.

## Chapter 2

### A Dogmatist Approach

#### 2.1. Introductory Remarks

Recall from the previous chapter that one central feature of emotional experience is that it instantiates phenomenal properties. In other words, just as there is something-that-it-is-like to taste coffee, or to see red and orange leaves on the ground, there is something-that-it-is-like to be indignant towards an unjust political policy, or proud upon achieving a significant goal, or distraught at the passing of a loved one, and so forth. The first account of emotional justification that this dissertation will consider identifies emotional phenomenology as the source of the emotions' ability to epistemically justify evaluative belief.

We can find support for this idea from various suggestive comments made by philosophers in the recent literature. Goldie (2004), for example, argues on behalf of an account of emotional experience “where the feelings involved are at center stage, playing a centrally important epistemic role in revealing things about the world” (p. 92). On a similar note, Tappolet (2016) argues that emotional experiences uniquely “allow us to be aware of certain features of the world” (p. 18), while Johnston (2001) claims that the epistemic import of affective experiences is rooted in their providing us with “affective disclosure” (p. 213) of evaluative properties. The focus of these claims on notions of ‘feelings’, ‘awareness’, and ‘affective disclosure’ certainly seems at least suggestive of the fact that these authors, among others, take the phenomenal properties instantiated by emotional experience – the *something-that-it-is-like* to undergo emotional experience – to bear epistemic significance.

In light of this support, then, how might we build a view of emotional justification based on the phenomenological character of emotional experience? One plausible way of doing so is to look to the wider philosophical domain of experiential justification; to look for a prima facie plausible account which bestows experiences with the power to immediately and defeasibly justify the relevant beliefs in virtue of bearing a specific kind of phenomenal character, and mould our theory of emotional justification on such an account. In this regard, one particularly plausible option is to appeal to a widely held internalist view of experiential

justification, namely, *phenomenal conservatism*.<sup>21</sup> According to phenomenal conservatism, an experience that makes it seem to you that *p* immediately and defeasibly justifies you in believing that *p*. On this account, if your visual experience makes it seem to you that there is an orange on the desk, then, absent defeating evidence, you have immediate justification for believing that there is an orange on the desk; if your intuitive experience makes it seem to you that it is morally wrong to kill one person to save five, then, absent defeating evidence, you have immediate justification for believing so.

Perhaps, then, we can draw up an emotional analogue of PC such that an emotional experience that makes it seem to you that *e* (where *e* signifies a proposition attributing an evaluative property to an object or event) immediately and defeasibly justifies you in believing that *e*. Call this view *emotional dogmatism*.<sup>22</sup> The purpose of this chapter is to begin an investigation into the prospects of an emotional dogmatist view by both: (i) clarifying emotional dogmatism and constructing the most plausible version of the view, and (ii) defending this version against a particularly conspicuous objection levelled against it in the literature.

The structure of this chapter is as follows. In §2.2, I begin by further elucidating the phenomenal conservative view and presenting the analogous emotional dogmatist thesis. In §2.2.1, I then argue that a liberal view which requires *only* that an experience bears a ‘seeming’ phenomenal character fails in virtue of being too permissive with respect to the kinds of experiences which end up getting ruled into the account. In §2.3, I then suggest that a more restrictive account of phenomenal conservatism, which requires something further in addition to this seeming character, is better placed to be a plausible model for an analogous emotional dogmatist view. Having set out restricted emotional dogmatism as the target view, §2.4 then presents a common objection levelled against the view, namely, the objection that evaluative properties are not suitable objects of emotional awareness. I will argue that this objection fails; that the restricted emotional dogmatist has two *prima facie* plausible stories to

---

<sup>21</sup> I appeal to phenomenal conservatism here, rather than phenomenal *dogmatism*, in light of the former constituting a broader model of experiential justification. Phenomenal dogmatism is typically discussed as referring to the immediate justificatory power of perceptual experience specifically, rather than also including experiences such as intuitive experiences, introspective experiences, and so forth. For the purposes of this chapter, however, I take it that we can use either phenomenal conservatism or phenomenal dogmatism to set up the possibility of an analogous emotional view.

<sup>22</sup> I borrow this label from Brogaard and Chudnoff (2016).

tell with respect to evaluative properties figuring into the contents of emotional awareness. Accordingly, §2.4.1 concentrates on explaining how evaluative properties can be *direct* objects of emotional awareness, while §2.4.2 sets out an account of how they can be *indirect* objects of emotional awareness. I then provide concluding remarks in §2.5.

## 2.2. Basic Emotional Dogmatism

Let us understand phenomenal conservatism as follows:

Phenomenal Conservatism (PC): if it seems to S that  $p$ , then, in the absence of defeaters, S thereby has at least some degree of justification for believing that  $p$ .  
(Huemer 2007, p. 30)

Now, there are a few things to note here. First of all, PC is an internalist view of epistemic justification insofar as it identifies factors internal to the agent (i.e. an agent's seeming states) as sole epistemic justifiers, as opposed to factors external to the agent, such as a belief's being produced by a reliable process. Secondly, identifying the source of an experience's justificatory power in its bearing the character of 'seeming to S that  $p$ ' is to identify it in the experience's phenomenal character, i.e. the something-that-it-is-like to undergo the experience. Thirdly, 'seemings' (or 'appearances') are typically taken to be non-doxastic propositional attitudes. That is, my visual experience can make it seem to me, say, that the Müller-Lyer lines are of different lengths, despite my belief that they are the same length. Cashing out exactly what these non-doxastic propositional attitudes amount to is difficult, and is typically best illustrated by drawing a distinction between two different types of experience. Compare, for example, my having a visual experience of an orange on the desk with my believing that there is an orange on the desk. Intuitively, so the thought goes, there is a significant phenomenal difference between my having the visual experience and my having the belief.

Different variants of PC diagnose this phenomenal difference in broadly similar ways. Huemer (2001), for example, describes visual experiences as instantiating the phenomenal property of 'forcefulness' (p. 77), insofar as in the visual experience of the orange, 'it seems to one that something satisfying the content of the experience actually exists, here and now' (p. 79). Moreover, Pryor (2004) claims that visual experience – like my visual experience of

the orange – ‘makes it feel as though, *by* enjoying that episode, you can *thereby just tell* that that content obtains’ (p. 357). Tucker (2010), on a similar note, argues that these visual experiences bear the phenomenal quality of ‘assertiveness’ (p. 530), where this is to be understood as the experience, in some sense, delivering something of an assurance that the content of the experience is true. What all three of these claims have in common is that they all allude to the apparent fact that my visual experience truly seems to recommend to me that there is an orange in a way that my merely believing that there is an orange doesn’t appear to do; the former *insists* that I accept the content of my experience as true, whereas the latter doesn’t appear to do so in the same way. However these variants choose to characterise this distinctive phenomenal character, this *seemingness*, the essential commonality of these accounts is that they all take it to be what bestows visual experience with the power to immediately justify belief.

Insofar as we’re interested in building an account of emotional justification on the basis of this view, we can extend PC over to emotional experience as follows:

Emotional Dogmatism (ED): if S’s emotional experience makes it seem to her that *e* (where *e* signifies a proposition which attributes an evaluative property to an object or event), then, in the absence of defeaters, S thereby has at least some degree of justification for believing that *e*.

On this view, just as my visual seeming experience of the blue mug can immediately and defeasibly justify my belief that there is a blue mug, my emotional seeming experience of awe towards a painting can immediately and defeasibly justify me in believing that the painting is awesome or beautiful.

This view of emotional justification is attractive for a number of reasons. First, note that general PC views are praised in virtue of their ability to provide us with a simple and intuitive explanation as to how we receive epistemic justification for our beliefs about the sensible world; we’re justified in believing what we do because of the way the world appears to us in our experience. Analogously, for emotional dogmatists, we’re justified in our evaluative beliefs about the world because of the way it appears to us in our emotional experience. Moreover, recall from Chapter 1 (§1.1) that an attractive feature of the general justificatory thesis of emotion was that, in conferring *immediate* justificatory power to emotional



experience, it was able to dispel sceptical worries concerning our access to justified evaluative belief. On a similar note, PC is thought to provide us with an antidote to pernicious sceptical worries pertaining to the justificatory status of our everyday beliefs about the sensible world given that epistemic justification comes at a low price for PCs. All that's required for epistemic justification is that our experiences bear the right sort of 'seeming' character. Insofar as ED is built on the foundations of PC, it seems as though it can provide an analogous remedy for sceptical worries pertaining to the justificatory status of our everyday evaluative beliefs, and thereby putatively strengthens the original attraction of the justificatory thesis of emotion.

### 2.2.1. Objection: An Over-Generalisation Problem

However, in spite of these advantages, there is an immediate worry with identifying an experience's justificatory power in its bearing the character of '*seeming to S that p*'. Namely, the theory appears to be too liberal with respect to the types of experience it bestows with immediate justificatory power; it lacks the theoretical tools to exclude experiences which we would *not* consider to be capable of conferring immediate justification from enjoying its epistemic yield. A popular way of presenting this challenge is in terms of the following example from Markie (2005):

Suppose that we are prospecting for gold. You have learned to identify a gold nugget on sight but I have no such knowledge. As the water washes out of my pan, we both look at a pebble, which is in fact a gold nugget. My desire to discover gold makes it seem to me as if the pebble is gold; your learned identification skills make it seem that way to you. According to [PC], the belief that it is gold has prima facie justification for both of us. (p. 356 – 357)

Markie's challenge constitutes a serious threat for PC. The possibility of rogue desires causing and manipulating the content of seemings, and thereby having an influence over which of our beliefs enjoy immediate justification, is worrisome for any theory which attributes such epistemic significance to seeming states. Indeed, consider an emotional case. To borrow an example from Brady (2013), suppose that I'm on the hiring committee for a new job, and upon interviewing a particular candidate, I find myself experiencing a negative emotion that makes it seem to me that this candidate is duplicitous or otherwise a bad fit for

the job (p. 87). It would be implausible to claim that this emotional experience *alone* is capable of immediately justifying my belief that the candidate is duplicitous on the basis of its bearing seeming phenomenal character. However, insofar as PC and ED *only* identify unqualified seemings as the relevant justification-conferring states, it appears as though the views do not have the resources to exclude cases like those above. It cannot be true that it is *only* in virtue of an experience bearing this ‘seeming’ character that it is capable of immediately and defeasibly justifying the relevant beliefs, or else we would have to concede that Markie’s gold prospector’s wishfully-produced seeming that the pebble is gold is afforded the same justifying role as the skill-produced seeming of the mineral expert, or that the suspicious interviewer’s belief is justified on the basis of their rogue emotional seeming. Thus, these over-generalisation cases are problematic for theories that invest positive epistemic status in seeming states alone.

The staunch defender of dogmatism may resist this conclusion, however. In response to cases like those above, it is always open to proponents of the view to bite the bullet and allow that, in virtue of their bearing the right kind of seeming character, experiences like these *are*, in fact, capable of immediately and defeasibly justifying belief. That is, defenders of PC and ED might be perfectly happy to concede that their theory generalises to experiences like that of Markie’s gold-pro prospector or the suspicious interviewer respectively, but deny that this is particularly problematic. It may be counterintuitive to those who are not naturally inclined to internalist views, but this isn’t a decisive objection insofar as the views can plausibly diagnose the intuitive oddness of these cases in other ways, e.g. by pointing to the fact that it is only *defeasible* and not *ultima facie* justification that is conferred by these experiences and that our intuitions are not sufficiently fine-grained to track the difference between the two, and to the fact that this justification is easily and often defeated, and so forth. Thus, PC and ED appear to have a relatively straightforward escape clause such that they can disarm worries concerning the apparent profligacy of the accounts.

This form of bullet-biting strikes me as implausible for two related reasons. To illustrate, consider the following case. Imagine a weak-willed agent who finds themselves living within a community of racists, all of whom harbour xenophobic beliefs towards those from a different ethnic background to themselves. Out of a strong desire to fit in with this group, our weak-willed agent actively engages with these xenophobic beliefs. She listens to racist propaganda, attends community events celebrating the lives and exploits of racist historical

figures, and so forth. Over time, she comes to adopt these beliefs herself, such that she forms a network of biases towards particular ethnic groups. As such, upon encountering a person that belongs to such a group, she habitually and systematically has the seeming that this person is acting suspiciously.

Here's the first reason that bullet-biting is undesirable. It strikes me as plausible to suggest that these xenophobic seemings are attributable to the agent herself and, specifically, to her desire to integrate into her community. She *created* and is responsible for the formation of those seemings. Now, ED, in virtue of its commitment to the claim that it is defeasible and not *ultima facie* justification conferred by experience, has the resources to explain why the agent's xenophobic seemings do not justify her in believing that the person from a particular ethnic group is acting suspiciously *only if* she has an awareness of her experience's etiology. That is, the justification conferred by the xenophobic seeming is defeated by her awareness of the fact that the seeming is ultimately attributable to her and her desires.<sup>23</sup> However, it *also* seems plausible that, as time passes and she successfully integrates into the community, she comes to forget the source of these xenophobic seemings, i.e. her desire to integrate into her community. Her racist beliefs become such an entrenched part of her cognitive architecture that she no longer questions them nor their origin. ED, then, appears to generate the strange result that the agent is *not* experientially justified in her belief that the person is suspicious at a time  $t^1$  where she is aware that her desire is the origin of the xenophobic experiential seemings, but she *is* justified on the basis of those seemings at a time  $t^2$  where she has forgotten that this is the case. This strikes me as counterintuitive. It is odd to suggest that *forgetting* something can enhance the positive epistemic status of a belief, *especially* when that belief is causally traceable and attributable to an agent's epistemically dubious desire. Thus, ED, even with its defeat clause, seems to be getting the wrong result here.<sup>24</sup>

---

<sup>23</sup> Huemer (2013), in defence of a phenomenal conservative view, argues that an experience's problematic etiology is epistemically relevant to the justificatory status of the relevant belief *only* when the subject has an awareness of that etiology.

<sup>24</sup> One might object that a case like this, i.e. a case in which an agent has *no* defeaters for the content of her xenophobic seemings, is unrealistic, given the many reasons that racist individuals typically have for doubting the accuracy of their xenophobic responses. In response, let me clarify that I don't intend for this counterexample to be representative of many real life cases, nor does it need to be in order for it to be successful. All that needs to obtain in order for the point to go through is that it meets the conditions for ED, and, yet, it seems counterintuitive to bestow justificatory power to the xenophobic seemings. For what it's worth, however, it's not obvious to me that this case is particularly contrived or beyond the realm of possibility.

Now, I take it that there are two ways in which the defender of ED might respond. First, the dogmatist may argue that, while there is *something* intuitively problematic about this case, it's not obvious that the problem pertains to the presence of epistemic justification. That is, one might contend that what our intuitions are *actually* tracking in this case is the agent's moral blameworthiness, or zetetic failings pertaining to her process of poor epistemic inquiry. If *these* failings are the source of our intuition that there is something amiss with this case, then the emotional dogmatist is let off the hook insofar as there's not *actually* anything problematic about bestowing her emotional seemings with justificatory power at  $t^2$ .

I take it that the best strategy for establishing that there *is* an epistemic failing here (and, specifically, one pertinent to the presence of epistemic justification) is to consider an analogous case in which there are no obvious moral or zetetic failings which plausibly hijack the intuition that there's something amiss with bestowing justificatory power to the emotional seemings. If we neutralise these non-justificatory failings and there's *still* something problematic about the epistemic result, then we have good reason to believe that this case *does* constitute an over-generalisation worry for ED. On that note, consider the following:

Wishful Affection: Suppose that, through a powerful desire to be liked by everybody, I come to believe that a person has strong affection for me whenever they remember my name. Consequently, I habitually experience the emotional seeming of joy whenever *anybody* refers to me by name; it emotionally seems to me that this referral is a very good thing for me. At a time  $t^1$ , when I am aware of these seemings' causal origin in my wishful thinking, they don't justify my evaluative belief that this event is good for me. At a later time  $t^2$ , when I have forgotten the etiology of these seemings, they *do* justify my evaluative belief.

Now, this Wishful Affection case shares the same general structure as the original over-generalisation case for ED. Plausibly, however, there's no obvious moral failing in this case. Moreover, it strikes me as unlikely that the issue at play is a zetetic worry pertaining to my poor process of epistemic inquiry given that I'm plausibly *not* conducting an inquiry when I have the emotional seeming of joy after somebody refers to me by name. According to Friedman (2019), a necessary condition for a subject to count as an inquirer, and to thereby have their process of inquiry subject to zetetic norms of assessment, is that they possess an "interrogative attitude" (p. 299) towards the question at hand, i.e. they're curious or

contemplative as to what the answer is. In this case, it's not obvious that I have the goal-directed activity of pursuing an answer to the question as to what any given individual's attitude is towards me; I just have the psychologically immediate experience of joy whenever a person refers to me by name, given my beliefs about what that referral means and my powerful desire to be liked. So, if a subject isn't morally or zetetically blameworthy in a case like this, but there *still* seems to be something counterintuitive about allowing their evaluative belief to be justified by their emotional seemings, then this seems best explained in terms of the subject's *specific* epistemic failing, such that bestowing their emotional seemings with justificatory power constitutes an over-generalisation problem for ED.

A second argument that the dogmatist might make in response to the over-generalisation case specifically concerns the worry that, for ED, forgetting key defeating evidence can improve the epistemic status of one's evaluative belief. To dispel this counterintuitive result, the dogmatist might appropriate argumentative resources from discussions of forgotten evidence and defeat in the epistemology of memory literature. One particularly relevant discussion concerns Huemer's (1999) proposal of the following diachronic view of phenomenal conservatism:

A belief is justified full stop if and only if one had an adequate justification for adopting it at some point, and thenceforward one was justified in retaining it. (p. 351).

This view is proposed partially in response to cases of forgotten defeat that are typically levelled against synchronic views of internalist justification. In these cases, a subject forms a belief that  $p$  via epistemically irrational means, such as wishful thinking. At a time  $t^1$ , when the subject is aware of this, her belief that  $p$  is unjustified. However, as time passes, the subject forgets the means through which she arrived at  $p$ , and retains  $p$  in memory at  $t^2$ . The worry is that many synchronic views will deliver the result that  $p$  is justified at  $t^2$  given that, at this time, the subject's defeater for  $p$  is lost to memory. Huemer's diachronic phenomenal conservatism attempts to avoid this result by claiming that a belief is overall justified if and only if the subject was once justified in adopting that belief, i.e., the subject's past mental states *matter* for the present justificatory status of one's belief. Given that, in the forgotten defeat case, the subject was never justified in adopting  $p$  because of its formation via irrational means, Huemer's view avoids the counterintuitive result. Returning to the case at hand, then, perhaps the emotional dogmatist can argue something similar. That is, assuming a

view like Huemer's, perhaps one can argue that the xenophobic subject is not justified in her evaluative belief that the person is acting suspiciously at  $t^2$  because the evaluative belief was not justified at  $t^1$ , given her then-awareness of her emotional seemings' etiology.

Here's the problem with this response. Even *if* diachronic views of this sort turn out to be plausible,<sup>25</sup> reasoning drawn from these discussions in the epistemology of memory cannot get a foothold on this over-generalisation case for ED given that, here, *nothing is being retained in memory*. Recall that, in the forgotten defeat cases pertinent to diachronic views like Huemer's, the subject forgets the defeating evidence but *retains* the belief that  $p$  via memory. The problem is that, in ED's over-generalisation case, the subject does not memorially retain the same belief that the person is acting suspiciously from  $t^1$  to  $t^2$ . Rather, at  $t^2$ , the subject has *another* emotional seeming experience which causes the belief which, crucially, is distinct from the belief formed at  $t^1$ . Because memory is playing *no* role here, plugging in a view like Huemer's will not be sufficient to dispel the counterintuitive result delivered by ED, nor can it absolve the dogmatist of the over-generalisation charge.

Thus far, then, we have a forgotten-defeater-based reason as to why bullet-biting in response to over-generalisation challenges is problematic. A second reason as to why such bullet-biting is worrisome pertains to the relationship between justified belief and rational action. One plausible principle prevalent in contemporary epistemology is that rational action is governed by some sort of epistemic norm.<sup>26</sup> That is, the question of whether an agent acts rationally is in some sense dependent on the quality of their epistemic position at the time of action. Suppose, for instance, that I have an important flight scheduled for 3pm and I spend the entirety of my morning lazing around, unconcerned by the approaching departure time. Bewildered by my nonchalance, my partner confronts me at 1pm and asks why I'm not already on my way to Edinburgh airport. I explain that it'll only take thirty minutes to drive to the airport, and a further fifteen to make my way through security, but they're unconvinced. "You don't *know* that it'll only take thirty minutes," they might tell me. Plausibly, they might stress that I can't know, nor am I justified in believing, that there won't be heavy traffic, or road diversions, or painfully slow queues at security. In this case,

---

<sup>25</sup> However, see Moon (2012a) and Smithies (2019) for arguments against diachronic moves of this type.

<sup>26</sup> Although, see Simion (2018) for the view that it is not an *epistemic* norm that governs action in general, but rather that it's more likely to be a moral or prudential norm with epistemic content.

intuition seems to suggest that I *shouldn't* be lazing around instead of making my way to Edinburgh because my epistemic base for doing so does not meet a sufficient standard.

While many endorse some instantiation of this unspecified principle, epistemologists disagree on what this norm actually *is*. Hawthorne and Stanley (2008) have suggested a *knowledge* norm of rational action, according to which you can “treat the proposition that p as a reason for acting only if you know that p”(p. 577). On this view, I can treat the proposition ‘it’ll take thirty minutes to drive to the airport and fifteen to get through security’ as a reason for continuing to laze around *only if* I know that this is the case. Since I cannot know that this is the case, I cannot rationally use this proposition as a reason for my action. Alternatively, some have elected to endorse a weaker epistemic standard, focusing on a norm of *justified belief* for rational action and practical reasoning, rather than knowledge.<sup>27</sup>

Now, for our purposes, all that matters is that there plausibly *is* some epistemic norm (or a non-epistemic norm with epistemic content) that governs rational action, and that the presence of epistemic justification *matters* for that norm to obtain. Returning to the xenophobe case, then, here’s the second reason that bullet-biting in response to over-generalisation cases is problematic. ED, insofar as it attributes justified belief to the weak-willed agent (when she lacks an awareness of the etiology of her experience), risks rationalising her actions on the basis of her xenophobic emotional seemings.<sup>28</sup> That is, assuming there is some sort of epistemic norm which governs action, and assuming that justified belief is at least necessary for this norm to obtain, ED does not only epistemically license harmful belief but *also* potentially harmful action. If the epistemic norm is justified belief, then ED straightforwardly puts the agent in a sufficient epistemic position to deploy the propositional content of her xenophobic experiential seemings in her practical reasoning. If the norm is knowledge, then although ED doesn’t stretch quite so far as to make it permissible for her to act, it still substantively contributes to her being in a sufficient epistemic position for action insofar as justified belief is putatively necessary for knowledge and thereby for the epistemic norm to obtain. The ease with which ED can rationally license (or at least contribute towards licensing) harmful action on such thin epistemic grounds strikes me as a serious mark against the view. Thus, insofar as biting the bullet and conceding

---

<sup>27</sup> See Littlejohn (2012) and Fantl and McGrath (2002) for different versions of this view.

<sup>28</sup> Brady (2013, p. 88) also gestures towards this point.

that ED easily generates justification to cases like the xenophobic seemings opens the door to normatively licensing vicious action (via epistemically licensing harmful belief), I take it to be a worrisome response.

Now, again, there are a couple of ways in which the dogmatist might respond to this argument. The first and most straightforward of which is to deny that there is any sort of epistemic norm which governs rational action. If a belief's being justified in no way influences whether it is rational to act on the basis of that belief, then there is no licensing problem for ED. This strikes me as a difficult argument to make. The dialectical weight in contemporary epistemology certainly seems to be in favour of the existence of such a norm, and, consequently, the defender of ED would have to both reject a commonly held view and explain away that intuition that some sort of epistemic warrant is required for rational action. While not an impossible task, it carries a significant theoretical cost.

A second means of responding to the objection may come from the dogmatist arguing that *even if* bestowing justification to a belief in some way positions that agent to rationally act on the basis of that belief, ED bestows only a small degree of justification. Recall from our initial formulation of PC that an experience's bearing the right sort of seeming character confers to the subject "at least *some degree* of justification for believing that *p*." Using this as our basis for ED, then, one might argue that the epistemic justification conferred by experiential (and, specifically, emotional) seemings to belief is *so* thin such that it provides only marginal rational support for acting. ED, therefore, does not substantively license harmful action, but rather certain actions enjoy only minimal rational support as a by-product of the view.

This also strikes me as an undesirable move. One reason for this is that in order for ED's rationalisation of harmful action to be intuitively non-problematic, the justification conferred by emotional seemings must be *so* minimal such that it raises the question of why it's worthwhile having in the first place. Presumably, one important desideratum for a plausible account of justification is that it delivers a robust enough notion of epistemic justification to belief such that we can then secure further epistemic goods, such as knowledge and understanding. Indeed, recall from Chapter 1 (§1.1) that one of the central motivations for endorsing the justificatory thesis of emotion is that it may constitute a naturalised account as to how we can receive further evaluative epistemic goods such as evaluative knowledge and



understanding, i.e. via the justification of our evaluative beliefs by emotional experience. If it is the case that ED can promise only a weak degree of justification, then it looks like a dogmatist approach to emotional justification will fail to satisfy this desideratum.

In summary, then, if identifying only unqualified emotional seemings as epistemic justifiers results in an overly permissive account of justification, and if endorsing such an account results in both counterintuitive and unsavoury implications, as I've argued here, then ED in its current form is not a plausible account of emotional justification. Emotional dogmatists must look elsewhere for a plausible epistemic framework for their view.

### 2.3 A Restricted Emotional Dogmatism

Now, returning our attention to PC, a fairly straightforward move for the defender of the view to make in response to cases like Markie's is to impose restrictions on the kinds of experiences that can immediately justify belief; to draw up a phenomenal conservative view which attempts to be more theoretically austere and precise so as to exclude the kinds of cases mentioned above. One example of a view particularly salient in the literature is that of Chudnoff's. Chudnoff, while taking it to be the case that certain experiences possess an epistemically significant phenomenal character, pushes the need for a more refined theoretical definition of this phenomenal character in order for us to properly investigate its epistemic implications. Chudnoff thus proposes that this phenomenal character is better understood in terms of what he calls *presentational phenomenology*:

Presentational Phenomenology: What it is for an experience of yours to have presentational phenomenology with respect to  $p$  is for it to both make it seem to you that  $p$  and make it seem to you as if this experience makes you aware of a truth-maker for  $p$ . (Chudnoff 2013, p. 37)

Here, Chudnoff sets out two conditions which must be met in order for an experience to have presentational phenomenology. The first condition Chudnoff specifies is plausibly understood as picking out the same phenomenal property gestured towards by Tucker and the other various PC accounts detailed in the previous section.<sup>29</sup> That is, an experience

---

<sup>29</sup> See Teng (2018, p. 642) for an argument in favour of this point.

possessing the property of making it seem to you that  $p$  is the same as an experience possessing the property of asserting or insisting to you that  $p$  obtains. The more interesting component of Chudnoff's analysis, then, is the truth-maker condition. On Chudnoff's view, it's not sufficient for my visual experience to make it seem to me, say, that there is a wasp on my sleeve in order for that experience to have presentational phenomenology with respect to that proposition, rather it must also be the case that my experience makes it seem to me as if I am visually aware of an item in my environment that corroborates that proposition, i.e. the wasp perched on the sleeve of my jumper. There are a couple of important things to note here.

First, Chudnoff identifies three central aspects of awareness in order to illuminate what exactly seeming-awareness of a truth-maker consists in. Suppose, for example, that I am aware of a small house spider on the wall. According to Chudnoff, it is a necessary condition that, in order for me to be aware of the spider, the spider at least partly 'contributes towards determining the phenomenal character of [my] experience' (Chudnoff 2018, p. 288). That is, the spider itself at least partly causally influences the *something-that-it-is-like* for me to undergo the visual experience of the spider. Moreover, it is a further condition of my having awareness of the spider that the phenomenology of my visual experience discriminates the spider from the range of background objects; I would not be aware of the spider if the phenomenology of my visual experience could not discriminate the spider from the other items in my visual field (if the spider was camouflaged amongst black polka-dot wallpaper, for instance). A final condition for my being aware of the spider is that I am able to 'entertain simple demonstrative thoughts' (p. 288) about the spider. Simply put, awareness of an object facilitates the ability to form thoughts which directly pick out the object *itself*, as opposed to picking out the object under certain descriptions or conceptualisations. While these are conditions imposed on general instances of awareness, they can be transposed to instances of *seeming* awareness insofar as all that is required for an instance of seeming awareness is that my experience *seems* to me as though it meets these conditions.

Second, note that presentational phenomenology is *relativised*. That is, it is not the case that certain experiences possess presentational phenomenology as a global phenomenal property consistent across all intentional contents, but rather that certain experiences possess presentational phenomenology localised to particular propositions. Suppose that you kindly buy me a cup of coffee and place it down on the table in front of me. Consequently, I have

the visual experience of a cup. I can see the shape, colour, and approximate volume of the cup, though I cannot see its contents. It is only on the basis of your testimony that I believe that the cup is filled with coffee. On Chudnoff's view, my visual experience has presentational phenomenology with respect to the proposition 'there is a cup on the table' (i.e. my experience both makes it seem to me that there is a cup on the table and makes it seem as if I am aware of a truth-maker for this proposition), but it does *not* have presentational phenomenology with respect to the proposition 'there is coffee in the cup'. This is because the contents of the cup are occluded from my vision; my visual experience thereby does not make it seem as if I am aware of a truth-maker for *this* proposition.

Chudnoff's phenomenological view then motivates a corresponding epistemological thesis, according to which it is only when an experience possesses presentational phenomenology with respect to  $p$  that it is capable of lending immediate and defeasible justification to the belief that  $p$ . This epistemological thesis will be referred to here as presentationalism:

Presentationalism: if an experience [immediately] justifies you in believing that  $p$ ... it does so in virtue of instantiating the property of having presentational phenomenology with respect to  $p$ . (Chudnoff 2012, p. 65)

On Chudnoff's view, it is *only* when my visual experience has presentational phenomenology with respect to, say, the proposition 'there is an orange on the table' that my experience can immediately and defeasibly justify my belief that there is an orange on the table. Presentationalism, then, can be understood as a version of PC restricted to experiences which instantiate the property of having presentational phenomenology.

There are many reasons to endorse this particular restricted phenomenal conservative account. One of the central motivations for accepting the view is that the notion of presentational phenomenology chimes well with various informal and metaphorical characterisations of the epistemically significant phenomenal character of perceptual experience offered by phenomenal conservatives in the literature, while providing a putatively more rigorous and theoretically robust diagnosis of this character.<sup>30</sup> For instance,

---

<sup>30</sup> Interestingly, appeal to a phenomenological view similar to presentationalism can also be found in the epistemology of memory. Namely, Hoerl (2001) distinguishes between propositional and episodic memory by arguing that, in instances of the latter, "the world as it was comes before the subject's

Chudnoff's truth-maker condition appears to map onto suggestions such as O'Shaughnessy's (2003) claim that 'whenever a person perceives-that p, he both believes that p and perceives something which is relevant to p's truth-value' (p. 319) or Tollhurst's (1998) claim that seemings are experienced as 'being revelatory of real features of the world' (p. 299). The focus here on perceiving an object or feature in the world pertinent to the truth value of the relevant proposition seems at least suggestive of something close to Chudnoff's characterisation of the truth-maker condition. Moreover, we can note that the presence of this condition determines that Chudnoff's account is better able to explain away over-generalisation cases. In Markie's gold prospector case, for example, Chudnoff can explain why my mere desiring that the pebble is gold is not sufficient for my experience to justify me in believing that the pebble is gold, i.e. my experience does not make it seem as if I am aware of a *truth-maker* for the proposition 'the pebble is gold' in light of the fact that my visual experience cannot make me aware of a truth-maker for the pebble's chemical composition.

Now, in light of this development, let's return to the emotions. We can transpose the theoretical machinery of presentational phenomenology over to the case of emotional experience in order to construct the following restricted account of emotional dogmatism:

Restricted Emotional Dogmatism (RED): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if the experience (i) makes it seem to her that *e*, and (ii) makes it seem as if she is emotionally aware of a truth-maker for *e*.<sup>31</sup>

On this view, it is not sufficient for my experience of awe to make it seem to me as though the painting is beautiful in order for it to have immediate justificatory power with respect to the relevant evaluative belief. Rather, my awe must also make it seem as if I am emotionally aware of a truth-maker with respect to the proposition 'the painting is beautiful'.

---

mind in such a way that it solves at the same time for what happened and for what puts her in a position to know what happened, namely that she was around to witness the event in question" (p. 332). Crucially, for Hoerl, this ability involved in the exercise of episodic memory is facilitated by a specific kind of phenomenological awareness. The emphasis on a memorial phenomenological awareness of the world as it was seems suggestive of something close to Chudnoff's truth-maker condition on presentationalism. Thanks to Fabrice Teroni for pointing my attention to this similarity.

<sup>31</sup> Cowan (2016) briefly discusses this view, characterising it instead as the emotions possessing "robust phenomenology" (p. 63).

One particularly interesting thing to note here is that RED's inclusion of the truth-maker condition fits nicely with the comments provided by Goldie, Tappolet, and Johnston in §2.1. Recall that in their respective descriptions of the epistemic power of emotions, Goldie described emotional feelings as being capable of "revealing things about the world", while Tappolet suggested that emotional experiences "allow us to be aware of certain features of the world". Again, the suggestion here that emotional experiences provide us with some sort of unique awareness about things *out there* in the world seems to closely match RED's requirement of emotional experiences making it seem as if we're aware of truth-makers for evaluative propositions, i.e. things *out there* that make evaluative propositions true. Indeed, regarding Johnston and his account of "affective disclosure", he *explicitly* uses the language of truth-makers insofar as he claims that "affect discloses evaluative truth-makers" (2001, p. 206), and that this (at least partially) explains what he terms the "epistemic authority" (p. 205) of affective experiences. In light of how well the comments of these notable authors fit with RED, I take the view to enjoy significant support from contemporary philosophers of emotion as well as from philosophers of perception and experiential justification.

To sum up, by including the truth-maker requirement, RED chimes well with views about the epistemic import of emotional phenomenology in the surrounding literature, inherits the general advantages of the basic ED account *and* receives support from a putatively more theoretically robust epistemological framework which avoids the pitfalls of PC. It is for these reasons that I will take RED as the most promising version of the emotional dogmatist view.<sup>32</sup> The remainder of this chapter will be concerned with defending RED against an objection commonly levelled against it in the literature.

#### 2.4 Objection: Evaluative Properties and Emotional Awareness

To reiterate, in order for my experience of awe to immediately and defeasibly justify my belief that a particular painting is beautiful, the experience must make it seem as if I am aware of a truth-maker for this proposition (in addition to making it seem to me that the painting is beautiful). Very plausibly, the truth-maker for this proposition is the evaluative

---

<sup>32</sup> I will, however, consider alternative instantiations of a restricted emotional dogmatist view in Chapter 3 (§3.3).

property of ‘beauty’ putatively instantiated by the painting.<sup>33</sup> In light of this, one common challenge levelled against emotional dogmatist views of this sort is that evaluative properties are *not* suitable objects of emotional awareness. If this is the case, then emotional experience cannot make it seem as if we are emotionally aware of truth-makers for evaluative propositions and, consequently, cannot bear immediate justificatory power according to RED.

Let us begin elucidating this objection by considering emotional seeming awareness. Recall that, in Chudnoff’s elucidation of presentational phenomenology (which constitutes the epistemological groundwork for RED), three conditions are specified for an agent’s seeming awareness of an object. That is, S has seeming awareness of an object *q* if and only if it seems to S that: (i) *q* at least partly contributes to determining the overall phenomenal character of S’s experience, (ii) S’s experience differentiates *q* from the range of background objects, and (iii) S’s experience enables S to form demonstrative thoughts about *q*.

Now, my *visual* experience of the painting plausibly meets the conditions imposed on instances of seeming awareness, i.e. the visual experience seems to differentiate the painting from background objects, it seems to enable me to entertain demonstrative thoughts about the painting, etc. It is therefore plausibly correct to say that my *visual experience* has presentational phenomenology, and is thereby capable of justifying my belief that there is a painting. But recall that presentational phenomenology is localised to particular propositions. That is, importantly, the fact that my visual experience has presentational phenomenology with respect to the proposition ‘there is a painting’ does *not* mean that it has presentational phenomenology with respect to the proposition ‘the painting is *beautiful*’. In order for my emotional experience to have presentational phenomenology with respect to *this* proposition, and thereby be capable of possessing the justificatory power we’re interested in here, it must make it seem as if I’m aware of a truth-maker for the evaluative proposition, i.e. the evaluative property of ‘beauty’ instantiated by the painting. Thus, my experience of awe must make it seem to me that: (i) the property ‘beauty’ at least partly contributes to the determining the overall phenomenal character of my awe, (ii) my awe differentiates the

---

<sup>33</sup> Brogaard and Chudnoff (2016) appear to assume that truth-makers for evaluative propositions are the evaluative properties putatively instantiated by the object or state of affairs in question. For the purposes of this chapter, this is an assumption that I will follow. However, as we’ll see in Chapter 3 (§3.2), questioning this assumption exposes a serious problem for RED.

beauty from the range of background objects, and (iii) my awe enables me to form demonstrative thoughts about the beauty of the painting.

How might my experience be capable of doing this? One way to understand it is analogous to the way in which visual experience introduces colour properties as objects of visual awareness.<sup>34</sup> Many have been tempted to draw a likeness between colour properties and evaluative properties on the basis of their both being plausibly identified as *response-dependent* properties. That is, just as the property of ‘redness’ cannot be understood without reference to the perceptual experience of observers, the property of ‘beauty’ cannot be understood without reference to the emotional experience of observers. In light of this similarity, then, it has been suggested that, just as visual experience can make us aware of colour properties, emotional experience can make us aware of evaluative properties.

The problem with this suggestion, and, indeed, the heart of the objection, is that it rests on an implausible account of what evaluative properties consist in. Colour properties are typically identified as *dispositional* properties, i.e. properties defined in terms of their being disposed to elicit certain responses in certain observers. ‘Redness’, for example, is typically defined as a property which is disposed to elicit visual experiences of redness among certain observers (i.e. observers in the absence of a malfunctioning sensory apparatus or bad lighting, for example). Evaluative properties, however, are *not* typically taken to be plausibly analogous to colour properties insofar as they are not plausibly understood as mere causal or dispositional properties of this sort. That is, an object’s being ‘beautiful’, say, is not plausibly understood as its being disposed to elicit responses of awe in certain observers in virtue of this view’s failing to account for the *normative* aspect of value. Rather, it is thought that an object’s instantiating the evaluative property of ‘being beautiful’ or ‘being amusing’ *merits* or makes fitting particular attitudes in response to the object, namely, awe and amusement respectively. Many take it to be implausible that the question of what is valuable can be distilled down to a brute psychological question concerning the attitudes that an object typically elicits in observers.

This difference between colour and evaluative properties leads us to the crux of the objection: evaluative properties, in virtue of their bearing this normative dimension, are not capable of

---

<sup>34</sup> See McDowell (1985) and Gert (2012)

being the objects of emotional awareness. My emotional experience of awe cannot make me aware (or even seem to make me aware) of the painting *meriting* responses of awe in virtue of the putative fact that whether an object merits or makes fitting a particular response is not something that I can be aware of via emotional feelings or phenomenology. This is the central objection facing RED in the literature. Brogaard and Chudnoff (2016) summarise the challenge as follows:

Suppose you fear a seen snake... You experience the snake. And you represent the snake as responsible for the bodily and mental turmoil. So far it seems you are just aware of truth-makers for propositions such as that there is a snake with a certain look moving a certain way, that you feel such and such unpleasant sensations and have such and such alarming thoughts. Are you also aware of an evaluative property of the snake, say fearsomeness or threateningness? Perhaps you attribute this property as part of the intentional content of your emotional response. But it is implausible that it is an object of awareness. The reason is that the property consists in the snake *meriting* and not just causing a certain range of bodily and mental changes in you. Whether the snake does merit those changes is not something that lies on the surface to be taken in by experience. (p. 70)

Here, Brogaard and Chudnoff identify this normative dimension of evaluative properties as the feature which bars them from being suitable objects of awareness. This idea is a familiar one in the philosophy of emotion. On a very similar note to Brogaard and Chudnoff, Dokic and Lemaire (2013) argue that it is “wildly implausible” (p. 236) to claim that emotional experience is capable of reflexively presenting itself (and its content) as fitting or unfitting. Relatedly, Deonna and Teroni (2012a) also assert that “When we undergo emotions, we seem to be entirely directed ‘outwards’ to the world and its properties, and in no way ‘inwards’ to responses that would be appropriate in the circumstances” (p. 101).

Summarising the discussion above, then, we can formulate the objector’s argument as follows:

- (1) We cannot be (seemingly) aware of normative properties via emotional feelings.
- (2) Evaluative properties are normative properties.



- (3) Therefore, we cannot be (seemingly) aware of evaluative properties via emotional feelings.
- (4) Evaluative properties are truth-makers for evaluative propositions.
- (5) As (3), we cannot be (seemingly) aware of evaluative properties via emotional feelings.
- (6) Therefore, we cannot be (seemingly) aware of truth-makers for evaluative propositions via emotional feelings.
- (7) If we cannot be (seemingly) aware of truth-makers for evaluative propositions via emotional feelings, then emotions cannot immediately and defeasibly justify evaluative beliefs.
- (8) Therefore, emotions cannot immediately and defeasibly justify evaluative beliefs.

Thus, the defender of RED faces a serious problem. In order for emotional experience to immediately and defeasibly justify one's evaluative belief, it must have presentational phenomenology with respect to the relevant evaluative proposition. Consequently, it must make it seem as if one is aware of a truth-maker for the evaluative proposition. However, insofar as the truth-makers for evaluative propositions are plausibly taken to be the evaluative properties instantiated by the relevant object or event, emotional experience *cannot* make it seem as if we are aware of these truth-makers in virtue of the fact that evaluative properties essentially involve a normative aspect. Therefore, in virtue of the nature of evaluative properties, and the corresponding limits of emotional awareness, RED fails.

In responding to this objection, and for the sake of argument, let us assume that evaluative properties are normative in the way assumed by the objector, i.e. that evaluative properties cannot be assimilated to mere dispositional or causal properties. The question now becomes whether defenders of RED have a plausible story to tell with respect to these evaluative properties being capable of being the objects of emotional awareness. I propose that they do. Or, more specifically, I propose that the objection paints too simplistic a picture with respect to awareness of evaluative properties. There are two crucial assumptions within the objector's argument: (i) evaluative properties cannot be direct objects of emotional awareness, and (ii) evaluative properties must be direct objects of emotional awareness in order for us to be emotionally aware of truth-makers for evaluative propositions. By evaluative properties being 'direct' objects of awareness, I mean the idea that evaluative properties can figure into emotional awareness in the same way that properties of shape figure into visual awareness.

As such, I take there to be two options open to the defender of RED. Firstly, corresponding to (i), they can argue that evaluative properties *can* be direct objects of awareness. Secondly, corresponding to (ii), they can argue that *even if* evaluative properties cannot be direct objects of awareness, they can be *indirect* objects of awareness, such that RED still has a story to tell with respect to us being aware of truth-makers for evaluative propositions. Here, I will elucidate each of these options in turn.

#### 2.4.1. Evaluative Properties as Direct Objects of Awareness

Now, as stated above, Brogaard and Chudnoff *and* Dokic and Lemaire take it to be straightforwardly true that emotions cannot make us aware of evaluative properties as properties which make fitting certain emotional responses. My fear of the snake cannot make me aware of the snake meriting or making appropriate emotional responses of fear, nor can it reflexively present itself as appropriate.<sup>35</sup> Indeed, recall that Dokic and Lemaire go as far as to describe the latter claim as “wildly implausible”. However, it’s not clear that this is true. That is, it’s not obvious that emotional experience can *never* contain self-referential content pertaining to its own epistemic appropriateness. If it *is* the case that emotional experience is capable of containing such content, then evaluative properties *can* be direct objects of emotional awareness, and emotional experience *can* make it seem as if we are aware of truth-makers for evaluative propositions.

For ease of explanation, call the kind of emotional content that we’re interested in *self-referential epistemic content*. We can locate support for the claim that emotional experience is capable of bearing such content from a range of sources. Take the metaethical literature, for example. As an early advocate for a version of the fitting attitude analysis of value, Brentano (1969) argues that our moral knowledge of what is good is ultimately derived from our having particular emotional experiences. That is, the good is that which we experience pro-attitudes towards which are *experienced as being correct*. These positive emotional attitudes are elucidated as follows:

When we ourselves experience such a love we notice not only that its object is loved and capable of being loved, and that its privation or contrary hated and capable of

---

<sup>35</sup> I will be using the notions of ‘epistemic fittingness’, ‘merit’, and ‘appropriateness’ interchangeably.

being hated, but also that the one is *worthy of love and the other worthy of hate*, and therefore that the one is good and the other bad. (p. 22, italics added)

Importantly, Brentano takes this normative element to be present in the experiential content of the emotion. That is, being aware of the emotional experience as correct and warranted by the object is a part of the emotional experience's phenomenological character. Alternatively, as a more contemporary example, Todd (2014) paints a picture of emotional phenomenology which builds in something like seeming awareness of fittingness conditions:

The claim I wish to defend is that although emotions do not have objective fittingness conditions, *it can appear to us that they do*, where this appearance of objectivity is subject to degree and best thought of in terms of a spectrum, ranging from the least objective-seeming to most objective-seeming responses. (p. 99, italics added)

Although Todd isn't explicitly clear about *how* we can be seemingly aware of these fittingness conditions in the experience of the emotion, it is clear that he takes it to be the phenomenological character of the experience that bears at least a substantial part of the fittingness content insofar as he takes it to be the case that "the relative fittingness of emotions will be more or less reflected in emotional phenomenology" (p. 102). From these examples, then, we can see that there is *some* support for the idea that the phenomenological character of emotional experience can bring awareness of its own epistemic fittingness.<sup>36</sup> In order to appease the objector, however, we need a *clear* account as to how emotional experiences can make us (seemingly) aware of themselves as epistemically merited with respect to their object. That is, we owe the objector an explanation as to *how* emotional experience can bear this normative dimension in its phenomenology, and why we have good reason to suppose that this is true.

One of the few specific accounts on the market as to how we can spell out this self-referential epistemic content can be found in Mitchell's work. Mitchell (2017) argues that emotions, in virtue of being valenced in a specific way, contain an awareness as to the "epistemic

---

<sup>36</sup> There are alternative accounts which I take to plausibly bestow (albeit implicit) support to this idea. See, for example, Montague (2009, p. 188)

appropriateness” (p.77) of the emotional response to object or event in question. Let me first summarise Mitchell’s view before clarifying it and defending it against objections.

Mitchell begins this line of thought with the claim that emotions possess what he terms “valence opacity” (p. 72). That is, emotional experience, unlike sense perceptual experience, is not phenomenally transparent insofar the phenomenal content of the experience is not exhausted by the object of the experience. When I reflect on my fear of the snake, I do not just reflect on the snake, but also on the emotional experience *itself*, or, more specifically, the affective component of the experience. My visual experience of a book, on the other hand, is phenomenally transparent. To put it in terms of a common metaphor, my visual perception constitutes a clear lens through which I ‘see’ the object or event in question. My visual experience *itself* brings no phenomenal features to the experience and I cannot attend to the visual experience itself. Emotions, however, constitute a lens tainted by valenced feelings of pleasantness or unpleasantness, approval or disapproval, favour or disfavour, etc.

From this claim, Mitchell then argues that these valenced attitudes that constitute an essential part of the emotion are typically felt *as appropriate* in response to the features of the object of the emotional experience. That is, I do not just experience my fear of the snake as involving felt disapproval towards the snake and its features, but I experience this felt disapproval as epistemically appropriate to or merited by the features of the snake, such as its venom, sharp fangs, etc. Mitchell elucidates this idea, and the central motivation for endorsing it, as follows:

So, the claim is that awareness of the epistemic appropriateness of one's felt valenced attitude, to what are experienced as content-external evaluative properties of the emotion's object [i.e. evaluative properties experienced as being possessed by the object *independently* of my having the particular emotional experience] is essential to those emotional experiences which are *intrinsically* rationally intelligible to us. Without experiencing the relevant value property, say danger, and in virtue of this experiencing one's felt disapproval as epistemically appropriate – as a merited uptake or registering of that value property – that fear would not be immediately rationally intelligible. (p. 78 – 79)

For Mitchell, then, emotional experience *must* contain an awareness of itself (or, more specifically, the felt valenced attitudes contained within the emotion) as epistemically appropriate with respect to the object of the emotion, otherwise we would be unable to make sense of our emotional experiences as immediately as we do. The claim, then, is that feelings of (dis)approval experienced as epistemically appropriate to the emotional object are *necessary* for the emotional experience to be immediately rationally intelligible to the observer. Moreover, given that many, if not most, of our emotional experiences *are* rationally intelligible to us, this seems to suggest that these valenced attitudes with self-referential content pertaining to the epistemic appropriateness of the emotion are fairly commonplace.

So, with this initial summary in place, why endorse Mitchell's valence-based account of self-referential epistemic content? Recall from earlier that the basic phenomenological thesis that emotional experience is capable of phenomenologically self-referring to its own epistemic appropriateness is not obviously a fringe view. We saw that Brentano and Todd made claims which support this thesis. Importantly, I take it we can *also* identify threads of support for the more specific idea that the emotions' opacity (compared to the transparency of sense perceptual experience) is at least indicative of their capacity for self-referential epistemic content from other authors. Cowan (2016), for example, makes a similar point, suggesting the following:

My suggestion is that the emotions' opacity may be thought to fit well with the suggestion that emotions possess NS content [i.e. self-referential content pertaining to their epistemic fittingness or appropriateness to the object of the emotion], i.e., if some emotions have NS content then it wouldn't be surprising that subjects refer, at least partly, to the emotion itself. (p. 25 – 26)

While Cowan doesn't make the stronger and more specific claim that Mitchell does, i.e. that the emotions' *valenced* opacity constitutes their self-referential epistemic content, it's clear that this suggestion maps onto the general idea behind Mitchell's phenomenological thesis. Moreover, as Mitchell notes, Poellner (2016) argues that emotions, insofar as they involve felt disapproval or approval to their objects, are experienced as "intelligibly motivated by the features of the object itself" (p. 266), which seems very close to what Mitchell has in mind. Thus, I take it that the fact that others have proposed views which share very similar philosophical roots to Mitchell's, *combined with* the general support available for the basic

phenomenological thesis, to be good evidence to believe that the objector in §2.4 moves too quickly insofar as they reject offhand the idea that there is any space in an emotional experience for self-referential content.

So, thus far, Mitchell's account looks like a potential candidate for a feasible answer to our objector. If we have good grounds for supposing that Mitchell's thesis is plausible, then we have good grounds for supposing that emotional experience is capable of possessing self-referential epistemic content, such that emotional experience *can* directly make us aware of evaluative properties, and thereby make it seem as if we're aware of truth-makers for evaluative propositions. Mitchell's view has a lot of moving parts, however, and requires some further elucidation in order to make it as plausible as it can be. In what follows, I will attempt to clarify Mitchell's account and defend it against some immediate objections.

First of all, there's a question of what it means to experience the felt valenced attitudes contained within the emotional experience *as appropriate* with respect to their object. For Mitchell, this is a particular kind of experienced causation. Crucially, Mitchell takes it to be the case that when I undergo an experience of fear towards the snake, I do not just experience the fear and then reflectively judge the fear to be caused by the snake, but rather I actually *experience* my fear as being seemingly caused by or as an apparent effect of the snake.<sup>37</sup> Now, of course, experiencing an emotion or felt valenced attitude as being merely caused by the object is not sufficient for that emotion or attitude to be experience *as appropriate* with respect to the object – further detail is needed here to bridge the normative gap.

While it's not *entirely* clear how appropriateness considerations get built into this experienced intentional causality, here's what I take Mitchell's explanation of this to be. Emotional experience, as we've seen, attributes an evaluative property to the object of the experience. My fear of the snake, for example, paints the snake in a particular evaluative light insofar as it attributes the property of 'fearsomeness' to the snake. Moreover, as Mitchell notes (and our objector will likely stress), normative reasons for experiencing certain emotions repose on evaluative properties.<sup>38</sup> The fearsomeness of the snake *provides normative reasons* for observers to experience fear towards the snake. Putting all of this

---

<sup>37</sup> Mitchell appears to take this to be an uncontroversial feature of emotional phenomenology; see (2017, p. 58)

<sup>38</sup> See Mitchell (2019, p. 128).

together, then, Mitchell takes it to be the case that in experiencing fear towards the snake, I both experience my fear as being caused by the snake *and* as responding to normative reasons in favour of experiencing fear, and *this* is what explains the fact that the felt valenced attitudes involved in the emotional experience are experienced as epistemically appropriate with respect to the object.<sup>39</sup>

As a second point for clarification, and relatedly to the previous point, recall that the central reason Mitchell cites for believing that these felt valenced attitudes are capable of being experienced as epistemically appropriate is that these experienced attitudes are necessary for emotional experiences to be *immediately rationally intelligible* to their subjects. Let us shed further light on this notion. What it means for an emotional experience to be immediately rationally intelligible is for that experience to non-doxastically ‘make sense’ to the subject in the moment that it occurs. That is, the experience is not perplexing for the subject insofar as the experience strikes her as rational or fitting with respect to the evaluative features of the object in a way that is *independent* of her having any post hoc reflective beliefs about the fittingness of the experience.

To illustrate, contrast the following two cases. In the first case, suppose that I, a seasoned presenter, have been scheduled to give an academic presentation to an audience of my peers, all of whom I’ve known for a very long time. Moreover, there are virtually no stakes in the matter; the worst that can happen is that I forget the material and face some light-hearted teasing in the pub afterwards. In spite of this, I find myself experiencing crippling bouts of nervousness several minutes before the presentation begins. I realise, however, that this anxiety is the result of my having drunk an entire pot of espresso an hour earlier. Now, in the second case, suppose that I, a very junior academic, have been scheduled to give a presentation to an audience of experts. The stakes are very high; the worst that can happen is that I forget the material and leave a very sour impression on people whose opinions of me matter for my future career prospects. Similarly to the first case, I find myself experiencing crippling bouts of nervousness several minutes before the presentation begins. I realise that this anxiety is the result of these high stakes.

---

<sup>39</sup> One might worry that this view has the implication that emotions are objectionably capable of becoming self-justifying. Mitchell appears to take the fact that the evaluative property attributed to the emotional object is experienced as being “content-external” is sufficient to allow him to sidestep the worry. See Mitchell (2017, p. 81 – 82) for this line of argument.

Now, in both cases, my experience of anxiety *makes sense* to me; my feelings of nervousness in both scenarios are intelligible and non-perplexing insofar as I have a satisfactory causal story which explains the experience. However, it is *only* in the second case that my anxiety enjoys immediate rational intelligibility in the sense that Mitchell has in mind. Recall from earlier that experiencing the felt valenced attitudes within the emotional experience as appropriate amounts to experiencing a particular kind of normatively-charged intentional causality with respect to the object of the experience, i.e. the felt valenced attitudes are experienced as appropriate with respect to the object insofar as they are experienced as both being caused by the object and responding to normative reasons in favour of experiencing those attitudes (and these normative reasons repose on the putative evaluative property attributed to the object by the emotion). Only in the second case do I experience my anxiety as making sense in terms of it being appropriately caused the object of the experience painted in a particular evaluative light, i.e. the high-stakes presentation. In the first case, while my anxiety is intelligible to me, it is not intelligible to me *because* I experience it as epistemically appropriate with respect to the object of the experience painted in a particular evaluative light, i.e. the low-stakes presentation.

At this point, it is important to note here that Mitchell's claim is *not* that all of our emotional experiences involve felt valenced attitudes experienced as epistemically appropriate to the object of the emotion, such that *all* of our emotional experiences enjoy immediate rational intelligibility. Mitchell invites us to consider instances of phobias as such cases in which these feelings of (dis)approval are *not* experienced as epistemically appropriate to the object, such that my emotional experience is rationally *unintelligible*:

Consider a certain type of arachnophobic whose felt disapproval (repulsion) towards a spider might not be experienced as merited by evaluative features of the object – as epistemically appropriate to it in this sense – but merely as habitually caused by it. What is partly characteristic of such cases is that the felt valenced attitude – and so the fear – is rationally unintelligible to the subject with respect to features of the object – hence we tend to describe them as emotional pathologies, and this is due to that feature of the experience. (2017, p. 78)



So my phobic fear of a harmless insect, for example, does not partly consist in an awareness of my valenced attitude, i.e. felt disapproval or disfavour, as being epistemically merited with respect to the features of the non-poisonous insect. Rather, the felt disapproval is experienced as only being *caused* by the object. Mitchell takes this experiential difference to be the reason why we tend to describe such as experiences as instances of phobic or pathological fear, rather than intelligible fear.

Here's one problem that might be raised in response to this claim. The objector may press the question of whether there is *really* an experiential or qualitative difference between phobic and intelligible fear, or, more broadly, between emotions that involve felt attitudes experienced as epistemically appropriate, and those that don't. Take another recalcitrant emotion case, for example. Suppose that, after having dreamed that my friend broke a promise to me, I experience anger towards them for the rest of the day. The anger is recalcitrant insofar as it both conflicts with and persists in the face of my opposing conscious evaluative judgment, namely, that my friend has not wronged me given that the transgression didn't actually happen. Now consider an instance in which, in reality, my friend *does* break a promise to me, such that I experience non-recalcitrant anger towards them for the rest of the day. Presumably, on Mitchell's account, the latter experience of anger involves a felt disfavour towards the object of my emotion (i.e. my friend's actions) which I experience as epistemically appropriate with respect to the evaluative properties of the object. On the other hand, Mitchell's view seems to rule that my recalcitrant anger involves felt disfavour which I *do not* experience as epistemically appropriate with respect to the emotion's object. One might argue, however, that there is no *felt* difference between the recalcitrant and the intelligible anger. In the moment that they occur, both experiences of anger have the same phenomenological character. If there is no qualitative difference between my recalcitrant anger and my intelligible anger, then this seems suggestive of the fact that, contrary to Mitchell's view, the experience of epistemic appropriateness is *not* a part of the emotion's phenomenology, nor intrinsic to the emotional experience itself.

In response to this objection, let's first look more closely at how an advocate of Mitchell's account might diagnose the recalcitrant case. Recall that, in ordinary cases, we experience the felt valenced attitudes contained within emotional experiences as epistemically appropriate insofar as we experience them as responding to normative reasons in favour of experiencing the given emotion, and these normative reasons arise from the evaluative property attributed

to the object by the emotional experience. Now, as we've seen, Mitchell suggests that in phobic cases (and presumably recalcitrant cases), the felt valenced attitudes are *not* experienced as epistemically appropriate with respect to the object of the emotion. One reason for this may be that the felt valenced attitude is not experienced as responding to normative reasons in favour of experiencing the emotion because the presence of the conflicting evaluative judgment constitutes a *defeater* for the normative pressure ordinarily exerted by evaluative properties. So, in fearing the harmless insect, my emotional experience attributes the property of 'fearsomeness' to the insect, but the felt disfavour is not experienced as responding to normative reasons in favour of experiencing fear because my evaluative judgment that the insect is *not* fearsome defeats those reasons conferred by the evaluative property putatively possessed by the insect. Put another way, the phobic or recalcitrant emotional experience instantiates intentional causality (i.e. is experienced as being caused by the object), but lacks the normative dimension to this intentional causality that is present in ordinary emotional experiences. Hence, the felt disfavour is not experienced as epistemically appropriate, and my fear does not enjoy immediate rational intelligibility.

Let me explain how this helps us in response to the objection. Certainly, there are *many* phenomenological similarities between recalcitrant (or phobic) emotional experiences and what we can call 'ordinary' emotional experiences for our purposes. Mitchell's account, as elucidated above, can explain this. That is, both recalcitrant and ordinary emotions are experienced as being caused by their objects, and as attributing evaluative properties to their objects. The *only* phenomenological dissimilarity between the two is that recalcitrant emotions are not experienced as responding to normative reasons in favour of experiencing that emotion. It seems plausible to me that this will not manifest itself as a *strikingly* salient experiential difference. Thus, Mitchell's account does not carve out significant and objectionable phenomenological space between these two experiences.

Having said that, the account does seem to entail a small experiential difference between recalcitrant and ordinary emotions insofar as, presumably, experiencing felt valenced attitudes (and thereby emotions) as appropriate and *not* experiencing them as appropriate are states with different qualitative flavours.<sup>40</sup> This does not strike me as problematic, nor as

---

<sup>40</sup> Perhaps there is an interesting difference between *not* experiencing a valenced attitude as appropriate and experiencing the attitude as *not* appropriate. That is, in the first case, the feeling of experienced appropriateness is merely absent whereas the second case claims the presence of a feeling

misdiagnosing the phenomenal character of recalcitrant and ordinary emotional experiences. Emotional experiences that we believe (or even know) to be missing the evaluative mark *do* seem to be experienced differently than those that we take ourselves to have no reason to doubt. The defeater that comes with a conflicting evaluative judgment seems to embed itself in the phenomenology of the recalcitrant emotional experience. There's nothing that *interferes* with ordinary emotional experiences in the same way. As such, because I take it to be plausible that there is a small phenomenological difference between recalcitrant and ordinary emotional experiences, and because Mitchell's account does not entail an implausibly *severe* experiential difference between the two (which I take the force of the objector's point to be), I do not think this objection spells significant trouble for Mitchell.

Now, with these notions of felt valenced attitudes experience as appropriate and immediate rational intelligibility in mind, why think that these felt valenced attitudes, experienced by the subject as epistemically appropriate to the object, are necessary for immediately rationally intelligible emotional experiences? One possible explanation for this is to return to the apparent fact that many of our emotional experiences are not mystifying to us in their occurrence. Typically, we're not perplexed by our emotional experiences, such that we don't feel the need to search for motivating reasons as to why we feel the way we do. Our emotions, in ordinary cases, *make sense* to us, and they do so *without* us having to engage in reflective reasoning after we've undergone the experience. Indeed, it would be psychologically implausible to suggest that we need to undergo the *post hoc* reflective reasoning in order for our emotional experiences to appear intelligible to us. So, if it's not the case that conscious cognitive mediation is what (typically) makes our emotional experiences rationally intelligible, and if many of our emotional experiences are rationally intelligible, the putative fact that emotional experiences themselves have this capacity *built in* insofar as they essentially involve felt valenced attitudes seems to be a good reason for thinking that these felt valenced attitudes are necessary for rational intelligibility.

A final objection one might level against the account at this point is that emotional experience bearing self-referential epistemic content *is not* the best explanation for the putative fact that most of our emotional experiences appear to us to be rationally intelligible.

---

of experienced inappropriateness. I will understand the first to be what Mitchell meant, given the word choice in the quotation above.

Indeed, one might argue that our emotional experiences are typically non-perplexing for us because of the wealth of experience we have with respect to typical emotion-eliciting scenarios.<sup>41</sup> That is, instead of requiring the presence of self-referential epistemic content in order to explain why our emotional experiences typically do not prompt us to search for motivating reasons as to why we feel what we do, we can appeal to the simpler explanation that we come to *learn* which emotions are appropriate by gradual acquaintance with particular formative emotional contexts, i.e. those which de Sousa (1987) terms ‘paradigm scenarios’. These scenarios, roughly, are typically experienced in our infancy and involve the subject experiencing instinctive responses (e.g. crying, smiling, etc.) towards particular objects and events in a given situation. The subject then comes to relate their experience to the particular features of the situation-type and thereby comes to grasp emotion concepts insofar as they associate their emotional experience with the eliciting scenario. Future judgments concerning the appropriateness of the relevant emotional experiences are then made on the basis of how closely the token situation at hand resembles the relevant paradigm scenario. Drawing on this, then, we can say that our emotional experiences are typically immediately rationally intelligible to us insofar as we take them to share the relevant features with their eliciting paradigm scenarios. There’s no need to posit self-referential epistemic content to explain rational intelligibility.

In response to this point, I take it that the relevance of paradigm scenarios to immediate rational intelligibility need not actually be an objection to the account. Indeed, I take it that there’s a plausible version of the paradigm scenario story that is entirely compatible with Mitchell’s view. Let us grant that we come to grasp emotion concepts (and concepts of emotional appropriateness) at least partly via acquaintance with paradigm scenarios. Moreover, recall that Mitchell takes the felt valenced attitudes of (dis)favour to be experienced as appropriate with respect to the object of the emotion insofar as they are experienced as responding to normative reasons which repose on the evaluative property attributed to the object by the emotion, and that this feeling of appropriateness is what gives some of our emotional experiences immediate rational intelligibility. It seems plausible that the influence of paradigm scenario acquaintance enters the picture at the level of property attribution, i.e. my emotional experience attributes the evaluative property of ‘fearsomeness’

---

<sup>41</sup> Elgin (2008) suggests something along these lines, claiming that our ability to readily distinguish between appropriate fears and phobias reflects a “sophisticated understanding of when and to what extent [our emotional experiences] are trustworthy” (p.9).

to the snake in virtue of the situation's resemblance to the paradigm scenario for fear, and I experience the felt disfavour towards the snake as appropriate insofar as it is experienced as responding to the normative reasons in favour of feeling fear that repose on the evaluative property of 'fearsomeness' attributed to the snake. My acquaintance with the paradigm scenario is thereby playing *some* role in my fear being immediately rationally intelligible to me, but it is still the phenomenology of the valenced attitude that is the central vehicle for the immediate intelligibility of the experience. The advantage of this is that it can accommodate the importance of formative paradigm scenarios of emotion and the influence of later association with such scenarios without having to rely on the presence of psychologically improbable cognitive mediation and inference in the emotional heat of the moment.<sup>42</sup>

In summary, then, Mitchell's proposal seems to be a fruitful resource from which the defender of RED can harvest a means of responding to the original objection. That is, the putative fact that emotional experiences can and do make us aware of the valenced attitudes contained within them as being epistemically appropriate with respect to the features of the emotion's object appears to provide evidence against both Brogaard and Chudnoff's and Dokic and Lemaire's brute claim that notions of appropriateness (or fittingness, or merit) cannot be incorporated into emotional awareness. Thus, if emotions *can* make us aware, or seemingly aware, of themselves as appropriate emotional responses with respect to certain objects, then the gap between emotional awareness and normative evaluative properties can be bridged, and the possibility of emotional experience making us aware of evaluative properties can salvage some plausibility. Importantly, my aim here has not been to provide the reader with irrefutable grounds for believing that Mitchell's account is true. For the purposes of answering the objector in §2.4, I have only endeavoured to show that it is *not* obvious that self-referential epistemic content cannot be built into the phenomenological character of emotional experience; that the defender of RED has the theoretical resources available to her to answer this objection by appealing to something *like* Mitchell's account.

---

<sup>42</sup> For what it's worth, I am very receptive to the possibility of learning mechanisms like those involved in de Sousa's paradigm scenarios account playing a significant role in the ability of emotional experience to immediately justify belief. I will not discuss this specific topic any further here, given my immediate focus on the phenomenological objection facing RED, but note that I will detail my own account of immediate emotional justification which appeals to learning mechanisms like these in Chapter 6 (§6.5).

It is entirely possible, of course, that the defender of RED is averse to advocating such a view. There may be some reason why emotional experience, in principle, cannot contain self-referential epistemic content. In what follows, I suggest an alternative route to answer the objector.

#### 2.4.2 Evaluative Properties as Indirect Objects of Emotional Awareness

Alternatively, the defender of RED can argue that we can be aware of evaluative properties indirectly insofar as we can be aware of evaluative properties *in virtue* of being aware of other properties. Let me explain.

Firstly, consider the following passage from Tappolet (2016):

It is quite right that evaluative properties are not dispositional properties. Being admirable, say, is not merely being such as to cause admiration. There are many things that cause admiration, but are not admirable. The question, then, is why emotions cannot inform us about evaluative properties understood in non-dispositional terms. After all, sensory experiences can inform us about non-dispositional, or primary, properties. We see shapes, for instance, and shapes can figure in the content of perception without requiring informational enrichment. So, if one allows that primary properties such as shapes can figure in perception in the absence of informational enrichment, there appears to be no reason to doubt that evaluative properties can do so as well. (p. 44)

Here, in response to Dokic and Lemaire, Tappolet points out that we ordinarily consider many non-dispositional properties, e.g. shape, size, etc., to be suitable objects of awareness. If, then, evaluative properties are non-dispositional properties, then we are at least warranted in considering the possibility that evaluative properties can be suitable objects of awareness.<sup>43</sup>

Now, consider the way in which our visual perception makes us aware of properties of shape, for instance. Plausibly, we can be visually aware of shape properties *in virtue* of being aware

---

<sup>43</sup> Soon after this passage, Tappolet explicitly considers the possibility that the normativity of evaluative properties bars them from being objects of awareness. She dismisses this on the grounds that advocating on behalf of such a view would amount to “simple-minded empiricism” (p. 45).

of colour properties. When I have the visual experience of a note-pad as being rectangular, for instance, I am visually aware of the note-pad instantiating this shape property on the basis of my being able to demarcate the note-pad from other objects in my visual field, and at least part of this ability consists in my visually experiencing various properties of colour, such as hue, tint, shade, and so forth. To put it another way, shape properties are accessible in my visual awareness because of my discriminatory abilities, and my discriminatory abilities essentially involve my being visually aware of dispositional properties such as colour. To generalise, then, there are cases in which awareness of non-dispositional properties at least partially depends on awareness of other properties.

The defender of RED may be able to transpose this line of thinking over to the case of evaluative properties as truth-makers for evaluative propositions. That is, if there are circumstances in which we can be aware of non-dispositional properties in virtue of being aware of other properties, then perhaps we can be aware of evaluative properties in virtue of being aware of non-evaluative properties. Reconsider the snake case, for example. It might be thought that analogously to the way in which I am aware of the shape properties in virtue of being aware of the dispositional colour properties, I am indirectly aware of the evaluative property of the snake, i.e. the snake's meriting responses of fear, in virtue of being emotionally aware of the non-evaluative properties which constitute the fearsome-making features of the snake, e.g. the piercing fangs, the quick aggressive movements, etc.

To make sense of this latter claim, the defender of RED may appeal to the idea that emotional experience can make us seemingly aware of an object's non-evaluative properties which *give rise* to the evaluative property, which is, in turn, the truth-maker for the relevant evaluative proposition. To put it in terms of Elgin's (2008) suggestion, for example, the defender of RED might point to the fact that emotional experience enables me to "discern a pattern in what I would otherwise take to be separate facts" (p. 44) in order to support the idea that emotional experience can make me aware of, say, the properties which constitute the fearsome-making features of my situation before the snake. Or, on a similar note, consider the following remarks about emotional phenomenology from D'Arms (2005):

It is in the nature of [emotional] experiences to present themselves as sensitivities to something outside them. And what they present themselves as sensitivities to is a fairly restricted feature of the situation: a socially significant personal inadequacy, or

a threat to one's safety, for instance. A little introspection makes it obvious, I think, that feelings of shame, fear and so on just aren't about the advisability, or morally permissibility, of feeling precisely that way. They are about a feature of the circumstance in virtue of which this is a fitting way to respond. (p. 10 – 11)

D'Arms' comments here seem to suggest that he takes emotional experiences to not *just* perform a functional role of highlighting evaluatively relevant patterns of non-evaluative properties, but also that emotions are *experienced* as being such indications or 'sensitivities'. Moreover, note that D'Arms takes it to be introspectively clear that emotional experiences make us aware of features in virtue of which it is *fitting* to have certain emotional experiences. This, combined with Elgin's support, appears to be good grounds for believing that emotional experiences, even if they cannot put us *directly* in touch with evaluative properties (i.e. properties which merit or making fitting certain emotional responses), *at least* uniquely aid in our awareness and apprehension of the pattern or conjunction of non-evaluative properties which give rise to the relevant evaluative property for the particular evaluative proposition in question. Therefore, in virtue of my emotional experience making me aware of the properties which constitute the fearsome-making features of the snake, the defender of RED may argue that my emotional experience indirectly makes me aware of the evaluative property of 'fearsomeness', such that it makes me aware of a truth-maker for the relevant evaluative proposition.

To summarise, §2.4.1 - §2.4.2 has been concerned with defending RED on the grounds that the original objection stated at the beginning of the section relies on too simplistic a view of emotional awareness and evaluative properties. I have argued that there are potential stories that RED can tell with respect to emotional experiences making us aware of evaluative properties which merit particular emotional responses. The argumentative burden shifts back to the objector. If they wish to push the objection, then they must rule out the explanations of evaluative properties figuring into emotional awareness that I've suggested here.

## 2.5. Concluding Remarks

The purpose of this chapter has been to construct and make tenable a version of emotional dogmatism which identifies the emotions' phenomenal character as the root of their ability to epistemically justify belief.



§2.1 began by highlighting the fact that various authors in the philosophy of emotion appear to attach epistemic import to the phenomenological character of emotional experience, i.e. the something-that-it-is-like for the subject to undergo the experience. I then suggested that, in light of this, a plausible contender for an attractive account of emotional justification is a view based on phenomenal conservatism (PC). From this, we began with emotional dogmatism (ED), according to which S's emotional experience immediately and defeasibly justifies her in believing *e* if and only if that experience makes it seem to her that *e*.

In §2.2, I set out both PC and ED in more detail, and then argued in §2.2.1 that these basic views ought to be rejected insofar as they problematically over-extend epistemic justification to beliefs which cannot intuitively enjoy positive justificatory status. I argued that it is implausible for the defender of these views to bite the bullet in response to these over-generalisation cases for two reasons. The first of which is that the PC generates counterintuitive results with respect to epistemic defeat and awareness of an experience's problematic etiology, while the second is that ED runs the risk of licensing harmful action via licensing harmful beliefs.

In light of these considerations, §2.3 then suggested that Chudnoff's restrictive account of PC may provide a better epistemological framework for our corresponding emotional dogmatist view. From Chudnoff's theory, I constructed an analogous restrictive emotional dogmatist view (RED) which affords emotional experience immediate and defeasible justificatory power with respect to an evaluative belief *e* insofar as it both makes it seem to one that *e* and makes it seem as if one is emotionally aware of a truth-maker for *e*.

In §2.4, I then presented a common challenge levelled against views of this sort, namely, the objection that evaluative properties, as truth-makers for evaluative beliefs, cannot be present in emotional awareness in virtue of containing a normative dimension. We cannot be emotionally aware of evaluative properties, so the objection goes, because emotional experience cannot present itself as being epistemically merited by the object in question. I argued that this objection can be countered by two responses. The first of which claims that emotional experience *can* contain self-referential epistemic content, and thereby bring direct awareness of evaluative properties (§2.4.1), while the second claims that emotional experience can *indirectly* make us aware of evaluative properties, such that we can tell a

plausible story about emotional awareness of evaluative properties as truth-makers for evaluative propositions.

In summary, RED looks to be in good stead. In the next chapter, however, I will show that, despite its success in response to this challenge, RED faces two further objections which cannot be overcome.

## Chapter 3

### The Problem with Restricted Emotional Dogmatism

#### 3.1 Introductory Remarks

Thus far, I have aimed to construct a restricted emotional dogmatist (RED) view on the basis of Chudnoff's account of immediate experiential justification, and I have defended it as a particularly plausible instantiation of an emotional dogmatist view. Recall that we're understanding RED as follows:

Restricted Emotional Dogmatism (RED): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if the experience: (i) makes it seem to her that *e*, and (ii) makes it seem as if she is emotionally aware of a truth-maker for *e*.

By building our account on Chudnoff's view – which I have argued is plausibly taken to be a precisification of many phenomenal conservative (PC) views found in the literature – we arrive at an emotional dogmatist account which both preserves the spirit and retains the theoretical advantages of liberal views, while putatively being sufficiently specific to rule out problematic cases. I have also argued that RED is capable of withstanding a familiar objection levelled against it in the literature, namely, the objection that evaluative properties, in virtue of bearing a normative dimension, are not suitable objects of emotional awareness.

However, in spite of RED's success up to this point, there are further, more serious challenges facing the defender of the account. One of the central aims of this chapter is to present what I take to be the most significant objection facing the view, i.e. the objection that there is no plausible way of spelling out what truth-makers for evaluative propositions consist in. After establishing this, and explaining why this failure gives us good cause to reject RED altogether, I then survey and reject alternative instantiations of a restricted emotional dogmatist view. In highlighting the severe shortcomings of these emotional dogmatist accounts, I thereby endeavour to build an inductive case for the conclusion that the phenomenal character of emotional experience cannot be what makes it capable of immediately and defeasibly justifying evaluative belief.

Accordingly, the structure of this chapter is as follows. In §3.2, I present a novel and powerful objection facing RED: namely, the objection that the inclusion of the truth-maker condition raises a troubling dilemma for the view: either RED identifies evaluative properties *themselves* as the truth-makers for evaluative propositions (§3.2.1), in which case the view will continue to over-generalise, or it identifies the relevant conjunction of non-evaluative properties as the truth-maker (§3.2.2), in which case the view ends up being objectionably restrictive. In light of the failure of RED, §3.3 then considers the plausibility of two alternative instantiations of restricted views, before concluding that neither of them are capable of faring any better than RED. I provide concluding remarks in §3.4 by reflecting on the vices of emotional dogmatist views in order to illuminate where we ought to look next for a plausible view of emotional justification.

### 3.2. Objection: The Dilemma of Evaluative Truth-Makers

Here, I argue that RED's inclusion of the truth-maker condition spells serious trouble for the view. Specifically, I argue that RED faces a dilemma as to what seeming awareness of truth-makers for evaluative propositions consists in. To illustrate, reconsider Brogaard and Chudnoff's example of an experience of fear towards an approaching snake. In order for that experience of fear to immediately and defeasibly justify the evaluative belief that the snake is fearsome, the emotional experience must both make it seem to you that the snake is fearsome *and* make it seem as if you're emotionally aware of a truth-maker for that evaluative proposition. The question that I'm interested in here is the question of what that truth-maker *is*.

Recall from the previous chapter that explicit analyses of RED, such as the one provided by Brogaard and Chudnoff, seem to assume that the relevant truth-maker for an evaluative proposition is the evaluative property putatively instantiated by the object of the emotional experience. However, RED, as it has been expressed thus far, is in fact silent as to whether the truth-maker consists in the evaluative property of *fearsomeness* itself, say, or whether it consists in the non-evaluative properties instantiated by the snake that give rise to the evaluative property of fearsomeness, i.e. the sharp fangs, the aggressive movements, and so forth. Call these 'the evaluative property reading' and 'the non-evaluative property reading' of the truth-maker condition respectively. The problem is that neither of these options looks promising for RED. In what follows, I will present each of these interpretive routes, and

argue that both render RED a problematic and unsatisfactory account as to how emotional experience can immediately and defeasibly justify evaluative belief.

Before doing so, however, let me first clarify how this objection both relates to and differs from the challenge discussed and rebutted in the previous chapter. Recall that the worry in §2.4 was that RED ought to be rejected on the grounds that it builds phenomenologically unrealistic contents into the scope of emotional seeming awareness. That is, assuming that the truth-makers for evaluative propositions are the evaluative properties instantiated by the object in question, and assuming that evaluative properties are normative properties which *merit* particular emotional responses, emotional experience cannot make it seem as if we're aware of truth-makers for evaluative propositions in light of the fact that emotional experience cannot bring seeming awareness of *itself* as being merited by the object in question. Put another way, this was a *phenomenological* objection facing RED.

I argued that this phenomenological objection was insufficiently powerful to undermine RED. That is, by appealing to a variety of views and explanatory stories as to how evaluative properties might directly or indirectly figure into the contents of emotional seeming awareness, I argued that the brute rejection of this possibility without due consideration of the prospects of these views should not be sufficient to compel us to reject RED out of hand. However, that is not to say that I think that conceiving of evaluative properties as the truth-makers for evaluative propositions is altogether non-problematic. Indeed, I think that the more promising route for the objector to take is to focus not on RED's putative commitment to controversial phenomenological assertions, but on its commitments to controversial *epistemological* results; commitments exposed by disambiguating the truth-maker condition.

My reason for this is twofold. First, as we've seen, whether one finds Brogaard and Chudnoff's challenge compelling relies on their having the intuition that emotional experience cannot bear a very specific kind of self-reflexive phenomenology. This doesn't strike me as a commonly held intuition. Second, and relatedly, it seems at least *prima facie* plausible that our intuitions have significantly more reliability and argumentative traction within the domain of epistemological theorising, given the frequency with which counterexamples are cited as compelling objections to epistemological views. Our intuitions when it comes to specific introspective phenomenological claims, on the other hand, are plausibly less widely-shared, less reliable, and less dialectically compelling. For these

reasons, the remainder of this section will pursue the forthcoming epistemological challenge against RED.

### 3.2.1. Truth-Makers as Evaluative Properties

With this in mind, then, let's begin with the evaluative property reading of RED's truth-maker condition, which can be spelled out as follows:

RED<sup>EP</sup>: S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if the experience both makes it seem to her that *e* and makes it seem as if she's emotionally aware of the evaluative property putatively instantiated by the object.

Now, recall that the inclusion of the truth-maker condition seems to suitably restrict phenomenal conservatism in the perceptual case (i.e. recall that Chudnoff's presentationalism seemed to be capable of sidestepping Markie's gold prospector over-generalisation case). The problem for RED<sup>EP</sup> is that it's not at all clear that this reading of the truth-maker condition restricts emotional dogmatism at all.

Reconsider the case of Brady's suspicious interviewer. In this scenario, I, as the chair of the hiring committee, get a 'bad feeling' about the interviewee and consequently have the experience of an emotional seeming that the interviewee is duplicitous. Recall from the previous chapter (§2.2.1) that, intuitively, this seems like a case in which we *shouldn't* allow the emotional experience to bear even defeasible justificatory power. Merely having an emotional seeming that a person is duplicitous doesn't seem sufficient to epistemically justify one's belief that they *are* in fact duplicitous. This was the problem for the basic ED view, and the motivation for including the truth-maker condition in our account as a means of restricting the net of emotional experiences capable of justifying evaluative belief.

The worry is that RED<sup>EP</sup> can't exclude the interviewer's emotional experience of suspicion because their experience satisfies both the seeming condition and the truth-maker condition. That is, insofar as the emotional experience *already* makes it seem to the interviewer that the candidate is duplicitous (and they're not aware of any reason to distrust this seeming), then plausibly their experience of suspicion *also* makes it seem to them that the candidate

instantiates the property of ‘duplicitousness’. The evaluative property reading of the truth-maker condition doesn’t seem to be adding *any* further requirement to emotional dogmatism, given that any emotional experience which satisfies the seeming condition will *also* satisfy the truth-maker condition. What else could it mean for an emotional experience to make it seem to you that the candidate is duplicitous, other than making it seem as if you’re aware of the evaluative property of ‘duplicitousness’ putatively instantiated by the candidate? Naturally, then, RED<sup>EP</sup> will continue to over-generalise to problematic cases precisely because, in practice, it’s no different to ED.

At this point, the defender of RED<sup>EP</sup> may argue that the case, as it stands, is under-described. In response to the worry, they might attempt to re-describe the case in order to explain and motivate the plausibility of conceding justification here. They may suggest, for instance, that my emotional experience is in actual fact picking up on subtle duplicitous-making features of the job candidate, i.e. that my having the emotional experience arises out of my perceiving certain mannerisms and micro-behaviours reliably indicative of untrustworthiness, such as avoiding the gaze of the interview panel, excessive talking, smirking, etc. Thus, picking up on the suggestion presented in the previous chapter on *indirect* emotional awareness of evaluative properties (§2.4.2), the defender of RED<sup>EP</sup> might argue that the emotional experience, insofar as it makes it seem as if I’m aware of the property ‘duplicitousness’ instantiated by the candidate, is making it seem as if I’m aware of a pattern of non-evaluative properties which constitute subtle duplicitous-making features of the candidate. If this is the case, then conceding immediate and defeasible justification on the basis of these emotional seemings doesn’t seem at all problematic – my emotional awareness in this case is tracking epistemically relevant features of the object, such that it seems perfectly plausible that the emotional experience is capable of conferring some degree of epistemic justification to my evaluative belief.

The problem with this response is that RED<sup>EP</sup> lacks the ability to distinguish between a case like *this*, i.e. a case in which the emotional seeming awareness of duplicitousness is caused by a seeming awareness of a pattern of duplicitous-making features of the candidate, and a case in which the emotional seeming awareness of ‘duplicitousness’ is caused by epistemically dubious cognitive biases, e.g. suppose that the candidate is a woman and the interviewer is unknowingly biased against women. The worry is that, insofar as the epistemically relevant emotional seemings – i.e. the seeming that the candidate is duplicitous

and the seeming awareness of the evaluative property ‘duplicitousness’ instantiated by the candidate – can be grounded in *either* of these causal explanations, RED<sup>EP</sup> doesn’t have the tools to differentiate the good and bad cases; both types of emotional seemings (i.e. those produced by epistemically legitimate means and those produced by epistemically illegitimate means) have the same justificatory power. This seems like a bad result.

To further support this point, contrast the following two cases. Suppose that Agent A navigates the entirety of her adult life with perfectly calibrated embarrassment responses, such that *every* experience of embarrassment she has picks up on the genuine embarrassing-making features of her situation. When she commits a social blunder in front of her colleagues, for example, her experience of embarrassment tracks the relevant embarrassment-making features of the situation, such as how her action or remark has transgressed a social norm, the obvious discomfort of those around her, the uncomfortable silence that follows her social buffoonery, and so forth. On the other hand, suppose that Agent B’s embarrassment responses are *wildly* unstable due to his extremely low self-esteem. He experiences embarrassment not only when he makes genuine social blunders, but when the breeze dishevels his hair, or when he’s wearing a different coloured shirt to the people around him, or when he gently coughs in a quiet room, and so forth. Moreover, let’s stipulate that B never second-guesses these embarrassment experiences; that he’s not aware of any reason to suppose that they’re missing the mark. His emotional experiences, clearly, are *not* picking up on genuine embarrassment-making features of these situations, and are instead being brought about by epistemically dubious background factors. I take it to be very intuitive that a plausible account of emotional justification will have the resources to explain why A’s emotional experiences of embarrassment are better epistemically placed than B’s.

The problem is that both of these embarrassment types – A’s legitimately produced embarrassment and B’s illegitimately produced embarrassment – will meet RED<sup>EP</sup>’s conditions. Just as A’s experience of embarrassment will make it seem that her social blunder is embarrassing *and* make it seem as though her situation instantiates the relevant evaluative property, B’s experience of embarrassment will make it seem to him that his wearing a different coloured shirt is embarrassing *and* make it seem as though the event instantiates the evaluative property of ‘being embarrassing’. The epistemically relevant emotional phenomenology is the same in both of these cases (i.e. RED<sup>EP</sup>’s conditions are met in both scenarios), but the intuitive verdict on whether the agent’s embarrassment has justificatory



power, I contend, is different between A and B. This result seems indicative of the fact that RED<sup>EP</sup> is failing to locate what's *actually* doing the justificatory work in emotional experience.

Now, one response open to the defender of RED<sup>EP</sup> here is to suggest that we *can* suitably distinguish between these cases by pointing to a phenomenological difference between those emotional seemings produced by epistemically legitimate and illegitimate causes, and identify *only* the former as meeting RED<sup>EP</sup>'s conditions on emotional justification. That is, the defender of RED<sup>EP</sup> might argue that only emotional experiences like A's embarrassment *actually* meet the truth-maker condition, and hence are capable of immediately and defeasibly justifying the relevant belief.

Accordingly, borrowing from the Mitchell-inspired line of argument presented in the previous chapter (§2.4.1), one might argue that the case in which the emotional seemings are not legitimately produced *does not* involve those seemings being experienced as epistemically appropriate with respect to the object. If we take emotional seeming awareness of evaluative properties (and consequently emotional seeming awareness of truth-makers for evaluative propositions) to necessarily require the emotional seemings to be experienced as epistemically appropriate or merited by the object of the emotion (in virtue of the normative dimension of evaluative properties), then, in cases like B's embarrassment experiences or the illegitimately produced suspicion of the interviewer, the emotional seemings will not make it seem as if the subjects are aware of the evaluative property putatively instantiated by the object, and thereby will not make it seem as if they're aware of a truth-maker for the evaluative proposition. Therefore, we have an explanation as to why RED<sup>EP</sup> confers immediate justificatory ability to the suspicious emotional seemings produced by legitimate observations of the duplicitous-making features of the candidate or A's experiences of embarrassment, but does not afford such power to emotional seemings which are not produced by these observations, i.e. because the emotion is experienced as immediately rationally intelligible in the former case, but not the latter, and this experiential quality is necessary for genuine seeming awareness of evaluative properties.

The problem with this response is that it's not at all obvious that emotional seemings produced by legitimate observations of the non-evaluative properties which give rise to the evaluative property are necessarily experienced as epistemically appropriate with respect to

their objects, and that emotional seemings which are not produced by such observations are *not* experienced as epistemically appropriate. That is, the immediate rational intelligibility of emotional experiences does not seem to be necessarily determined by the experience's having a proper etiology, i.e. being caused by such legitimate observations of the relevant non-evaluative properties. It seems quite possible that I can have an emotional experience which has an improper etiology, in this sense, which I experience as epistemically appropriate (and thereby immediately rationally intelligible) with respect to the object of my emotion. Suppose that, having drunk several more espressos in the morning than I'm used to, I later react with a burst of anger when my friend informs me that they'll be several minutes to our planned meeting. Now, given that I'm ordinarily an even-tempered person, it seems plausible that my emotional experience arises *not* out of legitimate observations about the non-evaluative properties of the situation which give rise to the evaluative property, namely, those which give rise to the 'offensiveness' of her tardiness, but instead arises out of my ingesting excessive quantities of a stimulant. However, it strikes me that I'll likely experience my anger as epistemically appropriate with respect to my friend's tardiness *despite* my emotional seemings not being grounded in the appropriate features of their object.

In light of this, I take there to be good reason to suppose that, similarly, the suspicious interviewer's emotional seemings can be experienced as epistemically appropriate *despite* not being grounded in or caused by legitimate observations of the duplicitous-making features of the candidate, or that B's emotional seemings can be experienced as epistemically appropriate *despite* not being grounded in or caused by legitimate observations of the embarrassing-making features of his situation. Hence, appeal to Mitchell's notion of immediate rational intelligibility is unable to provide RED<sup>EP</sup> with the resources to confer justificatory ability to legitimately produced emotional seemings but *not* to those produced illegitimately. So even if the defender of RED<sup>EP</sup> can mitigate the intuitive implausibility of conceding justification in the suspicious interviewer case by pointing to the fact that the emotional experience is tracking the relevant duplicitous-making features, this does nothing to assuage worries in cases where the experience is *not* picking up on these features.

So, if the defender of RED<sup>EP</sup> cannot differentiate between the good and bad cases, what about simply conceding over-generalisation to the problematic cases? The problem here is that this would mean that RED<sup>EP</sup> can no longer be considered a substantive improvement on basic ED. Recall from the previous chapter (§2.2.1) that I provided two reasons against freely

conceding justification to epistemically dubious experiences. The first is that reliance on an agent's awareness of epistemic defeat to explain the intuition that we shouldn't attribute full justification in these cases generates counterintuitive results, while the second is that freely conceding justification risks licensing harmful action via a knowledge or justification norm of rational action. Let's elaborate on these reasons by applying them to the emotional cases at hand. Suppose that we concede justification to the suspicious interviewer whose emotional experience of suspicion is *not* caused by legitimate observations of duplicitous-making features of the candidate, but by illegitimate cognitive bias. One plausible response the defender of RED<sup>EP</sup> might have here is the appeal to the notion of defeat. That is, they might argue that conceding justification is not so implausible given that it's only defeasible justification conferred by the experience. If the suspicious interviewer has an awareness of the fact that their suspicion is being generated by epistemically dubious means (i.e. by certain illegitimate biases), then the justification conferred by the emotional experience is defeated.

I suggested in the previous chapter that reliance on the defeater clause doesn't mitigate the intuitive implausibility of conceding justification in these cases and indeed generates counterintuitive results itself. That is, the defender of RED<sup>EP</sup> who intends to pursue this line of response is confronted with the odd result that the positive justificatory status of the suspicious interviewer's belief that the candidate is duplicitous switches on and off depending on whether the interviewer is currently entertaining an awareness of the problematic etiology of their experience. Moreover, on a similar note, it seems odd to think that we can easily come to have such an awareness of the etiology of our emotional experiences. Indeed, as we saw from Elgin in Chapter 1 (§1.1), we often take our emotional deliverances at face value. Without an *obvious* reason to distrust one's emotional experiences (e.g. a conflicting evaluative judgment, significant dissent from trustworthy observers, etc.), it seems inaccurate to suggest that the average person engages in reflective reasoning concerning the origin of all (or even most) of their emotions as they occur. As such, it seems plausible that, in cases like the suspicious interviewer's, it's *very* unlikely that the subject in question will have an awareness of the problematic etiology of their emotional experience, assuming that such awareness is accessible at all. If this is true, then by consequence it's also very unlikely that justification conferred by illegitimately produced emotional experiences will be defeated by the presence of such an awareness. Thus, the notion of defeat brought about by awareness of the emotion's etiology seems unable to diffuse the charge of problematic over-generalisation.

The second reason against freely conceding justification to experiences like the suspicious interviewer's, recall, pertained to the link between epistemic justification and rational action. I suggested that, in accordance with a widely held view (i.e. that there is a norm with epistemic content which governs rational action), freely bestowing justificatory ability to experience runs the risk of licensing harmful action. That is, if an agent's possessing epistemic justification for the relevant belief is *at least* necessary for the epistemic norm to obtain, and thereby for that agent to have rational grounds for the relevant action, then RED<sup>EP</sup>'s leniency with attributing justification spells trouble insofar as it at least substantially contributes towards licensing harmful action. Assuming the existence of such a norm (which, I've argued, is a plausible assumption), the fact that the suspicious interviewer is justified in their evaluative belief that the candidate is duplicitous *at the very least* contributes towards their having rational grounds for acting on that belief, i.e. dismissing the candidate as a bad fit for the job. This seems like an unpalatable result. So, all in all, if the defender of RED<sup>EP</sup> insists on conceding justification to problematic cases, then they must account for these objections.

### 3.2.2. Truth-Makers as Non-Evaluative Properties

So, if RED's continued vulnerability to the over-generalisation problem is the result of construing truth-makers for evaluative propositions as evaluative properties *themselves*, and if conceding justification in these cases is out of the picture, why not abandon this claim and insist instead that the truth-maker for evaluative propositions is the relevant set of non-evaluative properties that gives rise to the evaluative property? This is the non-evaluative property reading of RED's truth-maker condition, and can be spelled out as follows:

RED<sup>NEP</sup>: S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if the experience both makes it seem to her that *e* and makes it seem as if she's emotionally aware of the set of non-evaluative properties that, if instantiated, would give rise to the relevant evaluative property, and so make *e* true.

The attraction of this reading is that, unlike RED<sup>EP</sup>, it avoids obvious over-generalisation cases like the biased interviewer. Recall that, in this case, the interviewer's emotional

seeming awareness of the candidate's duplicitousness is caused by their illegitimate bias against women. This case would *not* meet the requirements of RED<sup>NEP</sup> precisely because the interviewer's emotional experience is *not* making it seem as if they're aware of the set of non-evaluative properties that would make the proposition 'the candidate is duplicitous' true. Rather, their experience is being triggered by the combination of their sexist bias and their perception of the candidate's gender. Clearly, mere seeming awareness of the candidate's gender does not amount to seeming awareness of the candidate instantiating particular non-evaluative properties which would make the proposition 'the candidate is duplicitous' true. Thus, RED<sup>NEP</sup> avoids the charge of over-generalisation because it *can* epistemically differentiate between the good case (i.e. the case in which the interviewer's emotional seemings of duplicitousness are caused by their perception of duplicitous-making non-evaluative features of the candidate), and the bad case (i.e. the case in which the interviewer's emotional seemings of duplicitousness are caused by their perception of the candidate's gender and their bias against women).

The problem, however, is that RED<sup>NEP</sup> is now too restrictive. If we identify these conjunctions of non-evaluative properties as truth-makers, then very few of our emotional experiences would be capable of bearing justificatory power. It seems that only very basic emotional experiences, like fear of a snake or disgust towards spoiled milk, for example, are reliably capable of bringing the required wide-ranging emotional seeming awareness of the relevant non-evaluative properties that would make the relevant evaluative proposition (e.g. 'the snake is fearsome', or 'the spoiled milk is disgusting') true. Emotional experiences which do not figure into this very basic category often don't bring awareness of the relevant non-evaluative properties.<sup>44</sup> Consider, for example, an emotional experience of awe towards a piece of artwork. It strikes me as plausible that such an experience does *not* bring full seeming awareness of the non-evaluative properties which would make the proposition 'that artwork is beautiful' true, e.g. the particular arrangement of brushstrokes, the particular blend of colours and hues, etc. Or, similarly, take an experience of amusement towards a particular state of affairs. It seems plausible that such an experience does not bring seeming awareness of the particular amusement-making non-evaluative properties, e.g. the relevant contextual

---

<sup>44</sup> Echeverri (2019) makes a similar point about the limits of awareness when it comes to the non-evaluative properties which constitute or give rise to the evaluative property in question, arguing that there are many *relational* properties which contribute towards an object exemplifying a given evaluative property, most of which we are "typically" (p. 550) not aware of.

factors which make the situation amusing. Despite the absence of such fine-grained seeming awareness in these cases, it seems entirely possible that emotional experiences of this sort are capable of providing a positive epistemic contribution to the status of the corresponding evaluative beliefs. Thus, robbing these emotions of immediate justificatory power on the basis of their not fulfilling the strict phenomenological requirements for RED<sup>NEP</sup> strikes me as bad news for the view.

Here, there are two possible responses available to the defender of RED<sup>NEP</sup>. The first of which is to concede that, understood this way, the view ends up being restrictive but deny that this is problematic. Indeed, the defender of RED<sup>NEP</sup> might stress that the lesson to be learned from the over-generalisation problem is that we *should* be casting a narrow net around the emotional experiences capable of bearing justificatory power. We want to rule *out* cases in which emotional seemings look like they're not grounded in epistemically legitimate observations of the relevant non-evaluative properties, and the best way of doing this is to impose strict constraints on what counts as emotional seeming awareness of truth-makers. If a consequence of this is that relatively complex emotional experiences which do *not* bring seeming awareness of the relevant non-evaluative properties end up getting ruled out of the account (insofar as they do not make it seem as if one is emotionally aware of a truth-maker for the relevant evaluative proposition), then so be it.

The worry with conceding epistemic austerity here, however, is that one plausible desideratum for a compelling version of a justificatory thesis of emotion is that it can account for how a *broad* catalogue of our evaluative beliefs can be justified by emotional experiences. If endorsing RED<sup>NEP</sup> means that we can *only* consider very basic emotional experiences as capable of bearing justificatory ability, then our dogmatist approach to emotional justification is failing to provide a satisfactory picture of the justificatory capacity of emotional experience insofar as JTE will no longer be able to secure the wide-ranging theoretical benefits identified in Chapter 1 (§1.1), i.e. JTE will only be able to explain how *some* of our evaluative beliefs come to be immediately justified, but not others; it will only be able to make sense of the trust we have in *some* emotional experiences, but not others.

Secondly, the objector might argue that in these scenarios – take the amusement case, for example – my emotional experience *is*, in fact, making it seem as if I'm aware of the relevant collection of non-evaluative properties which would make the event amusing, I just can't

articulate *exactly* what those properties are. One suggestion in support of this might be something like the following. When prompted, i.e. when asked ‘what’s so funny?’, I can gesture vaguely towards the features of the situation that make it amusing, such as the particular comment made, the context in which it was made, and so forth, even if I can’t express the amusing-making minutia. In other words, I’m not at a *complete* loss as to what it is about the situation that makes it amusing, and this is all that’s needed for evidence of emotional seeming awareness of the relevant conjunction of non-evaluative properties. Therefore, we *can* tell some story about having emotional seeming awareness of the relevant truth-maker in these cases, and RED<sup>NEP</sup> doesn’t end up being objectionably restrictive with respect to the kinds of emotional experiences it bestows with justificatory power.

The problem with this response is that further ambiguity in what emotional seeming awareness of truth-makers consists in raises difficult questions for RED<sup>NEP</sup>. If all that matters for emotional seeming awareness of truth-makers is that the experience makes the subject capable of *gesturing* towards the non-evaluative features of the object which would make the relevant evaluative proposition true, then it becomes less clear that RED<sup>NEP</sup> is able to rule out problematic cases. Take the suspicious interviewer whose emotional seemings that the candidate is duplicitous are caused by sexist bias. Plausibly, their emotional experience of suspicion will make them capable of saying *something* about what seems to make the candidate duplicitous (e.g. “there’s just *something* about their behaviour”), but this still seems insufficient for the interviewer to be justified in their belief that the candidate is duplicitous. Substantively relaxing the notion of awareness in order to let in cases where the emotional experience doesn’t make it seem as if one is aware of (i.e. able to identify) all of the relevant non-evaluative properties runs the risk of letting the epistemically illegitimate cases like biased suspicious interviewer in through the back door.

Moreover, note that relaxing the notion of awareness in this way is a significant departure from our original description of RED insofar as it no longer looks like the property of presentational phenomenology is being instantiated by emotional experiences. Recall from Chapter 2 (§2.3) that, for Chudnoff, a necessary condition for an experience instantiating presentational phenomenology is that the experience enables the subject to entertain *demonstrative thoughts* about the object of their seeming awareness. That is, a necessary condition for a subject to possess seeming awareness of an object is that the subject is able to form and entertain thoughts that pick out that object *itself*, as opposed to picking out the

object under certain descriptions or conceptualisations. If an emotional experience need only enable the subject to vaguely gesture towards the non-evaluative properties of the object which would make the proposition true, then it looks like the experience falls short of meeting this demonstrative thought condition and thereby falls short of qualifying as possessing presentational phenomenological character with respect to the evaluative proposition. RED<sup>NEP</sup> would therefore lose its footing in Chudnoff's epistemological groundwork. Given this result, I take this line of response to be an implausible defence against the argument.

In summary, RED is confronted with a troubling dilemma. Either we identify evaluative properties *themselves* as the truth-makers for evaluative proposition (RED<sup>EP</sup>), in which case the view continues to over-generalise to problematic cases, or we identify the relevant set of non-evaluative properties as truth-makers for evaluative propositions (RED<sup>NEP</sup>), in which case the view rules *out* emotional experiences which, plausibly, are capable of immediately justifying the relevant evaluative beliefs. If endorsing RED means that we must commit to either an objectionably profligate account of emotional justification or instead one which is objectionably austere, then RED does not provide a suitable framework for thinking about the immediate justificatory power of emotional experiences.

### 3.3. Alternative Restricted Views

Recall that in the previous chapter, I advanced RED as what I took to be a particularly promising instantiation of the emotional dogmatist view. This was because the addition of the truth-maker condition allowed RED to match onto suggestive descriptions of the epistemic importance of emotional phenomenology in the literature, while also providing an account of experiential justification that putatively both inherited the advantages of liberal views and painted a more nuanced analysis of the epistemically significant phenomenal character possessed by certain experiences. However, we've *also* seen that, insofar as RED attempts to qualify basic ED with the truth-maker condition, the view is rendered implausible.

One question the reader might have at this point is whether there exists an alternative instantiation of a restricted emotional dogmatist view. That is, if the addition of the Chudnoff-inspired truth-maker condition fails to make ED plausible, then perhaps we can look elsewhere for an alternative additional condition that can crystallise the view. If there is



such a condition, then the prospects of a dogmatist approach to an account of emotional justification are positive, even in spite of RED's failure. In what follows, I will consider two alternative suggestions for a restricted view of emotional dogmatism inspired by restricted phenomenal conservative accounts provided by McGrath and Markie, and argue that neither of these views can provide a plausible framework for an emotional dogmatist view.<sup>45</sup>

### 3.3.1 Receptive Seemings Emotional Dogmatism

Recall the gold prospector example which began our discussion of the over-generalisation problem back in §2.2.1. In this case, the expert prospector's perceptual seeming that the pebble is gold arises from their learned identification skills, while the wishful prospector's perceptual seeming that the pebble is gold arises not from their expertise but from their desire to discover gold. The problem for basic PC, of course, was that it was unable to account for the intuitive verdict that, while the expert may be immediately and defeasibly justified on the basis of their perceptual seemings, it is implausible that the wishful prospector's seeming has the same justificatory capacity.

In light of such counterexamples, McGrath (2013) aims to construct a restricted version of phenomenal conservatism which manages to exclude problematic cases while also striving to retain the initial attractions of basic views. On this note, McGrath suggests that what's going wrong in cases like the wishful prospector is that the perceptual seeming has what he refers to as a "quasi-inferential" (p. 228) basis, i.e. the wishful prospector's perceptual seeming that the pebble is gold does not arise *directly* from perception but instead arises via an inference-like transition or 'jump' from the base perceptual seeming that there is a yellowish pebble. The relationship between the seemings here is '*quasi-inferential*' insofar as exchanging the seemings with corresponding beliefs containing the same propositional contents would render the transition as an instance of inference between beliefs.

Now, for McGrath, it is only seemings which do *not* have such a quasi-inferential basis – i.e. what he calls '*receptive* seemings' – which are capable of providing immediate and

---

<sup>45</sup> Importantly, note that these accounts are more revisionary of PC than Chudnoff's view. Both McGrath and Markie impose restrictions on how justification-conferring seeming states must be generated, whereas Chudnoff only offers an extra phenomenological constraint on experiences capable of bearing justificatory power.

defeasible justification for the relevant belief; a quasi-inferential basis effectively bars a seeming from acting as a source of immediate or foundational epistemic justification. At *best*, seemings with a quasi-inferential basis might be capable of conferring mediate justification to the relevant belief, but *only* if it is an epistemically permissible quasi-inference, i.e. only if the content of the basis seeming adequately supports the content of the quasi-inferred seeming. Importantly, though, note that quasi-inferred seemings *cannot* generate immediate or foundational justification themselves, they can only transmit mediated justification to the relevant belief via an epistemically permissible quasi-inference from the base seeming.

So, applying these details to the example at hand, the wishful prospector has a receptive perceptual seeming that there is a yellowish pebble. On McGrath's account, the prospector would be immediately justified in believing that there is a yellowish pebble on the basis of this seeming. However, the prospector's desire to discover gold intervenes and produces a quasi-inferred perceptual seeming that the pebble is gold. Because *this* perceptual seeming is quasi-inferred from the base perceptual seeming that there is a yellowish pebble, it is *not* capable of immediately and defeasibly justifying the prospector's belief that the pebble is gold. Moreover, we can see that this quasi-inference taking place is not an epistemically legitimate one. The seeming with the content 'there is a yellowish pebble' does not sufficiently support the content of the quasi-inferred seeming, i.e. 'the pebble is a gold nugget'. Hence, the wishful prospector is in no way justified in their belief that the pebble is gold on the basis of their perceptual seemings. McGrath thus summarises this restricted view of PC as follows:

One is *prima facie* justified in believing P when one has a receptive seeming that P. In free enrichment cases [i.e. cases in which a seeming arises due to cognitive penetration], the output seeming isn't receptive; it isn't part of one's basic evidence or the grounds one fundamentally is "handed". It therefore cannot provide foundational justification. Nor can it provide non-foundational justification, because it depends on something very much like bad reasoning, viz., a bad "quasi-inference." (p. 244 – 245)

If this looks like a plausible view with respect to perceptual seemings, then perhaps we can construct an analogous emotional dogmatist view as follows:

Receptive Seemings Emotional Dogmatism (RSED): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief that an object O instantiates an evaluative property E if and only if (i) the experience makes it seem to her that O is E, and (ii) this seeming does not have a quasi-inferential basis.

RSED, then, attempts to qualify basic ED by excluding emotional seemings which are quasi-inferred from other seemings from enjoying immediate and defeasible justificatory power.

Now, to some degree, the question of whether RSED constitutes an improvement on RED hinges on whether RSED gives us the right result in emotional over-generalisation cases; whether it correctly diagnoses what's going wrong with the suspicious interviewer's emotional seeming, for example, and has the philosophical tools to exclude it from being capable of conferring justification. The problem is that, while the notion of receptivity may be plausible with respect to perceptual seemings and perceptual over-generalisation cases, it's not obvious that it translates particularly well to the emotional case. There's a question of whether *any* emotional seemings are receptive, and not quasi-inferred from other seemings, given that emotions rely on cognitive bases. That is, unlike perceptions, emotions depend upon base mental states such as perceptions, memories, beliefs, and so forth. I can't experience fear in response to the approaching snake without in some way perceiving the snake and its fearsome-making features. The same is not true of visually perceiving the snake; my visual experience of the snake does not presuppose a further mental state in the same way that my emotional experience does. In light of this fact, then, we might wonder how *any* emotional experience can involve a seeming that an object instantiates a particular evaluative property *without* that seeming being quasi-inferred from non-emotional seemings pertaining to the non-evaluative features of the object.

This is a problem because, if it is the case that all or most emotional seemings are quasi-inferred from the seemings of their cognitive bases (i.e. perceptual seemings, introspective seemings, etc.), it looks like RSED can't explain the intuitive epistemic difference between legitimately and illegitimately produced emotional seemings. Reconsider two versions of the suspicious interviewer case. In one scenario, the interviewer's emotional seeming that the candidate is duplicitous is caused by legitimate observations of duplicitous-making features of the candidate, whereas the other scenario involves the emotional seeming being caused by illegitimate background biases. For RSED, what has to be going wrong in the bad case is that

the interviewer's emotional seeming that the candidate is duplicitous is quasi-inferred from another seeming, and is thereby incapable of lending immediate justification to the evaluative belief that the candidate is duplicitous. But, as we've seen above, it looks like both the good *and* the bad case involve quasi-inferred emotional seemings. If *merely* being non-receptive makes a seeming incapable of conferring immediate justification, then RSED generates the same result for both the good and bad cases of suspicious interviewer.

In response to this point, the defender of RSED might argue that the view can still explain the intuitive difference in epistemic capacity between the emotional seemings involved in both cases. That is, they may point to the difference in epistemic quality in each quasi-inference as that which explains the intuition that the emotional seeming produced by legitimate observations is better epistemically placed than the seeming produced by illegitimate bias. Recall from above that, on McGrath's account, a quasi-inferential basis need not rob the seeming of *all* of its justificatory power. If it is a *good* quasi-inference, i.e. if the content of the base seeming adequately supports the content of the quasi-inferred seeming, then the quasi-inferred seeming can transmit mediate justification to the relevant belief. The defender of RSED might argue that in the *good* case, i.e. the case in which the emotional seeming that the candidate is duplicitous is quasi-inferred from the perceptual seeming which has as its content the relevant conjunction of duplicitous-making non-evaluative features of the candidate (i.e. their behaviours and mannerisms), the quasi-inference is legitimate insofar as the content of the base perceptual seeming adequately supports the content of the emotional seeming. We can see this by replacing the seeming experiences with beliefs with the same content. It seems plausible that the belief that the candidate is acting in particular ways (e.g. excessive talking, avoiding the gaze of the panel, etc.) is capable of lending inferential support to the belief that the candidate is duplicitous. On the other hand, consider the bias case. Presumably, the emotional seeming that the candidate is duplicitous will be quasi-inferred from perceptual seemings with *different* contents, e.g. if the bias is a bias against women, then the emotional seeming that the candidate is duplicitous will be quasi-inferred from the base perceptual seeming that the candidate is a woman. Clearly, this is not a legitimate quasi-inference; there is nothing about the base seeming that adequately supports the content of the quasi-inferred emotional seeming. In other words, there's an illegitimate 'jump' in the bias quasi-inference that isn't present in the good case, and this is what explains the difference in epistemic status between the two cases.

Now, even if this is a plausible way of explaining the intuitive difference between the two suspicious interviewer cases, it still doesn't get us where we want to go. Recall that we've been interested in how emotional phenomenology can *immediately* justify our evaluative beliefs. If it is the case that emotional seemings can only ever *transmit* mediate justification generated by perceptual (or memorial, or introspective) seemings, then RSED cannot account for emotional experience as a source of foundational or immediate epistemic justification.<sup>46</sup> Given that the justificatory thesis of emotion is wholly interested in the latter, endorsing RSED seems to concede too much. For these reasons, then, it looks like McGrath's receptivity-based view is not a suitable theoretical framework for a plausible restricted emotional dogmatist view. RSED does not fare any better than RED.

### 3.3.2 Knowledge-How Emotional Dogmatism

Finally, let's consider Markie's view. Returning to the gold prospector case, a natural suggestion as to why the expert prospector's perceptual seeming enjoys justificatory power is that the expert *knows* what gold looks like; the novice doesn't have anything close to this knowledge. One way of spelling out the problem with basic PC is that it can't account for the fact that this ought to make for a difference between the epistemic status of the expert's and novice's belief. In light of this natural intuition, Markie (2013) proposes a qualified view of phenomenal conservatism which restricts the type of seemings capable of possessing justificatory power to seeming experiences brought about by the agent's exercise of the relevant *knowledge-how* capacity:

The difference between them is that [Expert] knows how to visually identify gold nuggets, and his seeming state and resulting belief are an instance of his exercise of this know-how; [Novice]'s seeming experience is not an instance of any such know-how. Perhaps then an epistemically appropriate perceptual seeming experience to the effect that something is Q is one that is had in the exercise of the subject's knowledge of how to perceptually identify something as being Q. (p. 262)

---

<sup>46</sup> Although, see Cowan (2018) for the proposal that emotions of this sort (i.e. emotional experiences quasi-inferred from their cognitive bases), while incapable of providing *immediate* justification, can still serve as generative sources of justification for evaluative belief.

For Markie, then, merely having a perceptual seeming is insufficient for immediate and defeasible justification. A further condition must be met, namely that the subject must have the relevant knowledge-how capacity to recognise the relevant property *and* the seeming must be appropriately related to that capacity, i.e. the knowledge-how plays a substantive causal role in bringing about the seeming. Now, importantly, a subject's possessing the relevant knowledge-how capacity is not the same as their possessing the ability to reliably identify the property in question. Instead, on Markie's view, what possessing a knowledge-how capacity amounts to is the subject possessing a disposition to experience the relevant seemings upon perceptually apprehending certain features of the object in question, e.g. the expert prospector has the knowledge-how capacity to perceptually identify gold nuggets insofar as they are disposed to have the perceptual seeming that a pebble is gold when apprehending certain gold-making features of the object. Moreover, that subject's disposition is, as Markie puts it, "determined by" (p. 264) their having the right sort of background information, e.g. that an object which has certain features and looks a certain way is gold. Finally, on Markie's account, having this background information is a matter of having evidence that justifies the subject in believing, in this case, that an object which looks a certain way is gold.

So, if this view looks like its generating the right result in the perceptual case, we can transpose it into an emotional dogmatist view as follows:

Knowledge-How Emotional Dogmatism (KHED): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief that an object O instantiates an evaluative property E if and only if (i) the experience makes it seem to her that O is E, and (ii) S's experience makes it seem to her that O is E in virtue of her knowledge of how to emotionally identify something as being E.

To take a simple case, my emotional experience of fear towards the snake is capable of immediately and defeasibly justifying my evaluative belief that the snake is fearsomeness if and only if my experience makes it seem to me that the snake is fearsome, *and* I have this emotional seeming as the result of my knowledge of how to emotionally identify something as fearsome. Analogously to the details of the perceptual case above, a subject's knowledge-how capacity to emotionally identify something as fearsome (or disgusting, or enviable, or amusing, etc.) does not amount to mere reliable identification of fearsomeness, but rather

involves the possession of a disposition to experience emotional seemings of fearsomeness upon attending to certain features of the object or situation. Moreover, and again analogously to the perceptual case, I possess this disposition at least partly by virtue of my having the relevant background information, i.e. what *makes* fearsome things fearsome.

The good news for KHED is that it looks like the addition of the knowledge-how condition on seemings places the view in a much better position than RED to be able to handle over-generalisation cases. Take the case in which the interviewer's suspicious emotional seemings towards the candidate are produced by an epistemically illegitimate bias as opposed to epistemically legitimate observations of duplicitous-making features of the candidate. KHED seems to be able to provide a straightforward explanation as to why the interviewer's emotional seemings do not justify them in believing that the candidate is duplicitous, i.e. the suspicious seeming is experienced by virtue of the interviewer harbouring illicit biases, *not* by virtue of their knowledge of how to emotionally identify duplicitousness. The interviewer whose emotional seemings *do* arise as a result of legitimate observations of duplicitous-making features, however, plausibly does enjoy justification for their belief that the candidate is duplicitous insofar as their experiencing the suspicious seemings as a result of those legitimate observations are an exercise of her knowledge-how capacity to identify duplicitousness.

On a similar note, KHED is well-placed to explain the intuitive difference in justification attribution to A's perfectly-attuned and B's unstable embarrassment experiences. Plausibly, what explains the fact that A's embarrassment experiences consistently hit the evaluative mark is that she has developed a knowledge-how capacity to emotionally identify situations which instantiate the property of 'embarrassingness', and it is because of this that her embarrassment seemings confer justification to her relevant evaluative beliefs. B, on the other hand, clearly does not have this knowledge-how capacity, as evidenced by the instability of his embarrassment responses. KHED, then, seems to be in good stead insofar as it appears to have the theoretical resources to sidestep the over-generalisation problem.

However, I take it that there are two serious problems for the defender of KHED. The first worry is that KHED ends up being overly restrictive insofar as it over-intellectualises the acquisition of emotional justification. If it is the case that emotional experience can *only* justify evaluative belief when the experience is the result of exercising a knowledge-how

capacity to emotionally identify objects as instantiating the relevant evaluative property, and if possessing this knowledge-how capacity is determined by the subject possessing the relevant host of background information (i.e. evidence that justifies the subject's belief that an object instantiates a particular evaluative property), then it looks like emotional justification is going to be difficult to come by.

To illustrate, reconsider moral and aesthetic emotional experiences. There seem to be many instances of these experiences that we intuitively take to be good candidates for experiences capable of conferring justification to evaluative belief, e.g. indignation in response to discriminatory political policies, guilt upon breaking a promise to a friend, awe in response to a beautiful painting, and so forth. If KHED is true, then it is a necessary condition on these experiences being capable of conferring justification that they ultimately derive from the subject's having all of the relevant background information, i.e. evidence that would justify the subject in believing the relevant evaluative proposition, i.e. that the discriminatory policy is unjust, etc. This seems to be requiring too much of the subject. Moreover, it's not entirely clear that this account paints a plausible picture of *immediate* justification. Since justification-conferring emotional seemings must be the result of an exercise of a knowledge-how capacity, and since this capacity is determined by the possession of background information that would justify the relevant evaluative proposition, it's not at all obvious that KHED is capturing the phenomenon that we set out to explain, i.e. how emotional experience can confer justification to evaluative belief which does *not* depend on the subject's possessing any other justified beliefs. Like McGrath's account, then, KHED doesn't seem to be able to straightforwardly or adequately account for immediate emotional justification.

Relatedly, a second worry for KHED is that attributing so much weight to the possession of the relevant background information that determines one's disposition to have the relevant emotional seemings (and thereby the relevant knowledge-how capacity) threatens to render emotional phenomenology epistemically superfluous. That is, there's a serious question of what justificatory work the emotional seemings are doing if the brunt of the epistemic labour has already been done by the subject insofar as she putatively has the background information required to justify her belief that a given object instantiates the relevant evaluative property. Given the epistemic centrality and significance of phenomenal seemings for dogmatists, I take it that KHED's shift away from the epistemic importance of



phenomenology would constitute a mark against the view as a plausible version of emotional dogmatism.

To summarise, then, while KHED initially looks promising insofar as it *seems* better placed to handle the over-generalisation problem, we see on closer inspection that it both over-intellectualises the acquisition of emotional justification and, by the same token, threatens to make emotional phenomenology epistemically superfluous insofar as it places the brunt of the epistemic burden on the possession of background information about evaluative concepts. Moreover, because the view plausibly over-intellectualises the acquisition of emotional justification, I take it that KHED will not be attractive to those intuitively aligned with phenomenal conservatism. On the other hand, insofar as KHED *still* attributes epistemic weight to phenomenal character, I take it that it will also be unattractive to those *not* intuitively aligned with phenomenal conservative views. KHED, therefore, does not appear to fare any better than RED as a restricted emotional dogmatist view.<sup>47</sup>

### 3.4. Summing up

Thus far in the dissertation, my aim has been to investigate the prospects of a dogmatist approach to emotional justification, i.e. an approach which identifies the phenomenological profile of emotional experience as the source of its ability to lend immediate and defeasible epistemic justification to evaluative belief. I have endeavoured to show that there appears to be no plausible way of spelling out the epistemically significant phenomenal character of emotional experience and, as such, that there appears to be no plausible candidate for a dogmatist approach to emotional justification.

To recap, the beginning of this chapter saw RED enter the discussion with a clean slate. In the previous chapter, I identified RED as the most promising instantiation of the emotional dogmatist view, and defended it against the objection that evaluative properties, in virtue of bearing a normative dimension, are incapable of being objects of emotional seeming

---

<sup>47</sup> While KHED strikes me as an implausible view for the reasons mentioned here, note that my own view, which will be elucidated in Chapter 6 (§6.5), shares certain similarities to KHED. As will become clear, I think that mechanisms of emotional learning are central to the emotions' ability to perform justificatory roles. These learning mechanisms *could* be considered similar to Markie's construal of the knowledge-how capacity, but, crucially, I do not take learned emotional capacities to consist in the possession of articulable background information.

awareness. However, in §3.2, I argued that the question of what truth-makers for evaluative propositions consists in leads RED to a troubling dilemma. §3.2.1 saw that if we take evaluative properties to be the relevant truth-makers ( $RED^{EP}$ ), then the view will continue to over-generalise, whereas §3.2.2 established that if we instead take the relevant set of non-evaluative properties to be the truth-makers ( $RED^{NEP}$ ), then the view will under-generalise and fail to include emotional experiences which, intuitively, are capable of lending epistemic justification to evaluative belief. In light of RED's shortcomings, §3.4 then considered two alternative instantiations of the restricted ED view: Receptive Seemings Emotional Dogmatism (RSED) and Knowledge-How Emotional Dogmatism (KHED). I argued that neither of these views fare better than RED.

So, where does this discussion leave us? The failure of both basic and restricted emotional dogmatist accounts, it seems, is that emotional *awareness* is not a sufficiently sophisticated philosophical tool through which we can secure a satisfactory account of emotional justification. First, our scope of emotional awareness was cast too wide, in which case *too many* emotional experiences were bestowed with justificatory power, but it was also seen that our attempts to adjust and fine-tune the scope of emotional awareness by imposing further restrictions on what can be plausibly included resulted in much too narrow a focus. The fact that emotional dogmatist accounts continually miss the mark when it comes to our intuitive verdict on whether an emotional experience can justify evaluative belief seems suggestive of the fact that emotional awareness and phenomenology just *isn't* what's doing the justificatory work; that there's something else that's determining whether emotional experiences are capable of justifying evaluative belief beyond the what-it-is-like-ness to undergo the experience. The prospects for a dogmatist approach to emotional justification, then, are bleak. I propose that we ought to look elsewhere for a plausible version of the justificatory thesis of emotion.

## Chapter 4

### A Reliabilist Approach

#### 4.1. Introductory Remarks

Thus far, we've seen that the epistemic capacity of emotional experience to immediately and defeasibly justify evaluative belief cannot be plausibly spelled out in terms of emotional experience possessing an epistemically significant phenomenal character. Indeed, we saw that determining the justificatory status of a given evaluative belief on the basis of whether the subject's emotional experience bears the relevant seeming phenomenal character repeatedly generated counterintuitive and problematic results. One possible explanation for this is that the question of whether a given emotional experience instantiates a particular phenomenal character is not in *any* way bound up with whether that emotional experience accurately represents its object. That is, focusing solely on a subject's emotional awareness and her internal seeming states neglects the question as to whether our emotional experiences are *actually* accurately tracking the presence of evaluative properties, and *this* might seem significant with respect to determining whether emotional experience can immediately and defeasibly justify evaluative belief.

In light of this, then, it seems like a natural transition following our rejection of phenomenology-based views of emotional justification is to shift our focus from *internalist* aspects of emotional experience, i.e. emotional awareness and phenomenology, and move to an account which invests epistemic significance in *externalist* aspects, i.e. factors external to the subject's mental states. We can locate threads of support for an externalist view of emotional justification from a number of authors in the literature. Specifically, many of these authors unite around the notion of emotional *reliability* as the epistemically relevant externalist factor. For example, Elgin (2008) hinges the putative justificatory power of emotion on their being "reliably correlated" (p. 38) with the way things are in the evaluative landscape. Additionally, Pelser (2011) defends a version of the justificatory thesis plausibly understood as reliabilist insofar as he seemingly allows that the emotions' pervasive and widespread unreliability would be sufficient to render them epistemically powerless.<sup>48</sup>

---

<sup>48</sup> Strictly speaking, Pelser takes the justificatory thesis to be compatible with a wide range of both internalist and externalist theories of justification. However, due to his extensive defence of the thesis against the unreliability problem, I take it to be a plausible interpretation of his position that he

Moreover, Brady (2013) appears to assume a reliabilist conception of the justificatory thesis in his discussion of the view, claiming that “[the] epistemic role and value of emotional experience... depends upon the existence of reliable causal links between emotional experience and the occurrence of “core relational themes” of danger, insult, contamination, less, shamefulness, wrongness, and the like” (p. 76).<sup>49</sup> From these comments, then, we can draw up our next contender for an explanation as to how emotional experience can immediately and defeasibly justify evaluative belief: *emotional reliabilism*. Emotional reliabilism, roughly, is the view that emotional experience is capable of lending immediate and defeasible justification to evaluative belief in virtue of the fact that basing one’s evaluative belief on the relevant emotional experience is a reliable process of belief-formation.<sup>50</sup> In this chapter, I aim to do the following: (i) motivate and clarify the most plausible version of the emotional reliabilist view, and (ii) defend emotional reliabilism from the most conspicuous objection facing it, namely, the unreliability problem.

This structure of this chapter is as follows. In §4.2, I set out emotional reliabilism in more detail by explaining the motivations for adopting it, and highlighting the philosophical fruit that doing so might bear. I also clarify the view by rebutting three preliminary objections facing emotional reliabilism in §4.2.1, §4.2.2, and §4.2.3 respectively in order to pave the way for the central critical discussion of the chapter. §4.3 then begins this discussion by presenting the unreliability problem and explaining why it ought to concern emotional reliabilists. I then identify what I take to be the three most promising ways in which one might argue for the systematic unreliability of emotional experience and defend emotional reliabilism by rejecting all three as implausibly theoretically costly. I then provide concluding remarks in section §4.4.

## 4.2. Motivating and Clarifying Emotional Reliabilism

Emotional reliabilism is based on a process reliabilist view of epistemic justification.

Following Goldman (2008), let us understand process reliabilism (henceforth ‘reliabilism’) in terms of the following:

---

attributes a great deal of epistemic importance to the reliability of emotion. See Pelsner (2011) for further detail.

<sup>49</sup> Note again, however, that Brady rejects the justificatory thesis of emotion as implausible.

<sup>50</sup> In this chapter, and throughout the dissertation, note that I often use ‘emotional reliability’ as shorthand for the reliability of forming true evaluative beliefs on the basis of emotional experience.

Reliabilism: a proposition  $p$  is propositionally justified for an epistemic agent  $S$  (at time  $t$ ) just in case  $S$ 's total mental state (at  $t$ ) is such that *if*  $S$  were to “target” proposition  $p$  and were to apply suitable reliable processes in her repertoire to that total state, then a belief that  $p$  would be generated in  $S$ . (p. 77)

We can unpack this with a straightforward case of sense perceptual experience. Take the visual experience of an orange on the table. For reliabilists, I'm immediately and defeasibly justified in believing that there's an orange on the table if and only if that belief is formed via a reliable cognitive process (or processes). In this case, the relevant process involves my visually experiencing the orange and forming the belief that there is an orange on the basis of that visual experience. Now, because that belief-forming process is reliable insofar as it tends to produce more true beliefs than not, I am thereby immediately and defeasibly justified in believing that there is an orange on the basis of my visual experience. My merely *guessing* that there is an orange, on the other hand, would not be capable of justifying my belief that there is an orange, given that guessing lacks a favourable ratio of true output-beliefs over false ones, and thereby fails to constitute a reliable belief-forming process.

Reliabilism is typically considered to be the chief representative and paradigm of externalist epistemology, and enjoys a host of support from epistemologists in the literature.<sup>51</sup> The primary attraction of the account is that it secures an intuitive and robust connection between justification and truth. To illustrate, reconsider the case of Markie's gold prospector.

Reliabilists can easily explain why the novice prospector's belief that the pebble is gold fails to enjoy the same positive justificatory status as the expert's, i.e. the wishful novice lacks any reliable processes which would give rise to the belief that the pebble is gold. Indeed, recall that the wishful prospector forms their belief on the basis of their desire to discover gold. Because this is evidently *not* a reliable process of belief formation, i.e. it will only by *chance* generate a true belief and will regularly produce false beliefs, the novice prospector is not justified in believing that the pebble is gold. The expert prospector, on the other hand, *is* justified insofar as their believing that the pebble is gold on the basis of their perceptual experience and learned identification skills *is* a reliable process.

---

<sup>51</sup> Reliabilist epistemology has been defended by many notable authors, including Goldberg (2010), Kornblith (2009), and Lyons (2009).

So, if reliabilism looks like it's generating the right result in the perceptual case, then we can build an emotional analogue as follows:

Emotional Reliabilism (ER): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if forming evaluative beliefs on the basis of emotional experience is a reliable process of belief-formation.

For ER, my experience of fear towards the snake is capable of immediately and defeasibly justifying my evaluative belief that the snake is fearsome if and only if my evaluative beliefs formed on the basis of my emotional experiences tend to be true. Crucially, for the emotional reliabilist, the reliability of the overall belief-forming process requires that the subject's emotional system reliably generates true evaluative content, i.e. the subject's emotional system tends to be reliably responsive to the presence of evaluative properties. Of course, reliability does not entail infallibility. It's perfectly compatible with ER that my emotional experiences occasionally miss the mark under unideal conditions. The only requirement is that forming evaluative beliefs on the basis of emotional experience is a reliable *enough* belief-forming process. As a final clarificatory point, note that, for the purposes of this dissertation, I do not intend to construe the truth of ER as a logical consequence of the truth of generic process reliabilism. That is, the prospects of ER will be evaluated independently of the truth or falsity of generic reliabilism.<sup>52</sup>

There are good reasons to endorse ER. The first of which, as we've seen above in §4.1, is that the view enjoys a host of support from notable authors in the literature. Moreover, one particularly attractive feature of this view from our perspective is that, similarly to the way in which general reliabilism offers a solution to Markie's prospector case, ER promises to provide us with a way out of the emotional over-generalisation problem that stumped dogmatist views. Recall the suspicious interviewer. ER has the resources to be able to explain the intuitive verdict that the interviewer is significantly more likely be justified in their evaluative belief when the emotional experience of suspicion is based on their observations of duplicitous-making features of the candidate, but *not* when the interviewer's suspicion is produced by illegitimate background biases. That is, very plausibly, the relevant belief-

---

<sup>52</sup> This dialectical point will be revisited in Chapter 5 (§5.2).

forming processes are different. Forming an evaluative belief on the basis of emotional experience *generated by* illicit biases is likely to be a largely unreliable process of belief formation, whereas forming evaluative beliefs on the basis of legitimate observations of the non-evaluative properties which give rise to the evaluative properties will be significantly more reliable. ERs can point to this difference in the reliability of the processes in each of the cases as what explains our intuition that the epistemic status of the latter case is better than that of the former.<sup>53</sup>

Now, in spite of these attractions, ER faces three preliminary objections. The first of which is a causal worry, the second is a version of a generality problem, and the third targets the necessity of reliability for epistemic justification. In what remains of this section, I will tackle these worries, defending and clarifying ER along the way.

#### 4.2.1 A Causal Problem

First, one might worry that ER suffers a *causal* problem. More specifically, it might be objected that there doesn't seem to be the right kind of causal nexus between emotional experience and evaluative properties required for a tenable reliabilist theory. Reliabilism, as a theoretical descendent of a causal process theory of justification, is traditionally understood to be a paradigm of a naturalised epistemology, locating epistemic value and epistemic properties in natural causal processes that take place between the believer and the objects of her belief. To illustrate, consider sense perceptual experiences. It might be thought that the reliability of my visual perception in tracking properties such as colour, shape, and motion is grounded in the particular causal interaction taking place between those properties and my perceptual faculty. That is, my visual experience of the book on the table is caused by the book's properties, my auditory experience of a passing car is caused by the car, and so forth. Thus, I am justified in believing that there is a book on the table or that a car has passed by at least partially in virtue of the fact that I stand in a certain causal relation to the objects of those beliefs.

---

<sup>53</sup> Here, one important question may arise concerning *how* we individuate the relevant processes which are to be evaluated in terms of their reliability. I address this worry in §4.2.2.

In order for ER to be appropriately analogous to general reliabilism, then, it might be argued that a similar causal interaction must take place between emotional experience and evaluative properties; that emotional experience and evaluative properties must be appropriately causally networked in order for subjects to have epistemic access to evaluative properties, and thereby for a plausible reliabilist conception of the justificatory thesis of emotion to get off the ground. However, it doesn't seem obvious that my emotional experience of fear is caused by the evaluative property of fearsomeness, or that my emotional experience of anger is caused by the evaluative property of offensiveness. Indeed, claiming that such evaluative properties *are* causally efficacious is a controversial metaethical position.<sup>54</sup> As such, it might be objected that forming evaluative beliefs on the basis of emotional experience cannot be a reliable belief-forming process precisely because our emotions are not causally 'hooked up' to evaluative properties in the right kind of way, and therefore, that emotional experience cannot be supplemented into a reliabilist theory of justification.

It seems to me that the defender of ER can dispose of this objection in one of two ways. First of all, one might deny that the kind of causal interaction detailed above is necessary for reliabilism.<sup>55</sup> As it has been spelled out here, all that ER requires is that the following precondition is met: forming evaluative belief on the basis of emotional experience is a reliable process of belief formation. It is not obvious that this precondition necessitates that emotional experience and evaluative properties are causally networked. That is, it's not obvious that the relevant causal interaction is required for reliability. Consider mathematical knowledge and justified belief, for example. It doesn't seem obviously incorrect to suggest that we reliably form true mathematical beliefs (e.g. that 12 is the square root of 144, that 3 is a prime number, etc.) without being appropriately causally networked to the relevant abstract mathematical entities. Thus, the defender of ER can diffuse the objection by detaching reliabilist theories of justification from the causal constraints imposed upon them by the theoretical context in which they arose.

However, for the sake of argument, suppose that tenable reliabilist accounts do require causal interaction between the believer and the object of her belief. This does not undermine ER, for the defender of the thesis *can* appeal to a certain causal connection between emotional

---

<sup>54</sup> However, see Oddie (2005) for a defence of this view.

<sup>55</sup> Shafer-Landau (2003) argues for this point in his response to a very similar causal objection facing his process-reliabilist theory of moral knowledge.



experience and evaluative properties. While it might be true that evaluative properties themselves are causally inert, one might insist that evaluative properties supervene on the non-evaluative properties of objects which *are* appropriately causally networked to emotional experience. On this account, while my fear may not be caused by the ‘fearsomeness’ property of the dog, it is plausibly caused by the non-evaluative features of the dog that the evaluative property supervenes on (e.g. the sharp teeth, the snarling maw, etc.). We can therefore understand evaluative properties as having *some* sort of causal power, albeit a causal power not intrinsic to them but derived from the non-evaluative, descriptive properties which constitute their supervenience base.<sup>56</sup> Thus, in light of these feasible responses, I take it that the defender of ER has the resources available to her to disarm causal worries of this sort.

#### 4.2.2 A Generality Problem

The second problem facing ER is a familiar foe for any instantiation of a reliabilist account. The generality problem for reliabilism, roughly, is the problem of non-arbitrarily individuating the processes of belief-formation which are to be assessed in terms of their reliability. Reconsider the visual experience of the orange. The token process of belief formation in this case, i.e. forming a belief that there is an orange on the basis of the visual perceptual experience of an orange, is an instance of many types. Not only is it a token instance of the type ‘visually perceiving an orange’, but also of ‘visually perceiving an orange in good lighting conditions’, ‘visually perceiving an orange on a Tuesday’, ‘visually perceiving an orange when I’m in a good mood’, and so forth. Now, the process *token* itself is a one-off and unrepeatable event, and cannot be assessed in terms of its reliability. It seems, then, that we must look to the process-*type* to which that token belongs as the locus of our reliability assessment. The question for reliabilists is *which one* of these process-types is relevant for the assessment of reliability. This is an important question, particularly given that it’s entirely possible that reliability verdicts, and thereby attribution of justification, will differ across the range of processes. The challenge for reliabilism, then, is to provide a principled way of determining which process-type is the correct candidate for reliability

---

<sup>56</sup> Audi (2013) endorses a similar causal story with respect to moral properties. On Audi’s account, moral properties are grounded in non-moral properties, such that the former can figure into causal relations but only in virtue of the causal power of the non-moral properties in which they are grounded.

evaluations. Without such a rule, reliabilism is, as Conee and Feldman (1998) put it, “radically incomplete” (p. 3).

With this in mind, we can see how an objector might level a generality-style challenge against ER. That is, one might raise the challenge of how we are to identify the relevant process type for forming evaluative beliefs on the basis of emotional experience (e.g. ‘fear in response to slithery animals’, ‘fear in response to a snake’, ‘fear in response to a snake which is grounded in observations of the fearsome-making features of the snake’, etc.). This is a question concerning how fine the process-grain ought to be when it comes to reliability evaluations.<sup>57</sup> It’s clear that this question is important given that we’ve already seen the impact the answer has on reliability verdicts. Recall from above that ER’s answer to emotional over-generalisation cases, like the suspicious interviewer, is to point to differences in the relevant process-type. That is, the interviewer’s evaluative belief *is* justified when it is based on an emotional experience which is in turn based on observations of the duplicitous-making features of the candidate (given the reliability of this belief-forming process), and *isn’t* justified when it is based on an emotional experience caused by cognitive biases (given the comparable *unreliability* of this belief-forming process). But on what principle can we carve a distinction between these processes, such that *these* are the relevant belief-forming processes to be assessed in terms of their reliability?

There are a couple of things to say about this. First of all, given that this particular problem for ER is an instance of a more *general* problem for reliabilists, I take it that ERs have a range of argumentative options to choose from in how they answer this challenge. Many reliabilists have attempted to provide a principle for type-individuating the relevant belief-forming process, choosing to individuate on the basis of such features as how the belief is based on the subject’s available evidence, the information-processing problems solved by the belief-forming process, the algorithms and parameters employed by the process, and so forth.<sup>58</sup> Given that ER strikes me as no more worse off with respect to the generality problem than these other general reliabilist views, and because these general reliabilist views have this range of argumentative options to choose from, I take it that the defender of ER can also make use of whichever option she deems most plausible.

---

<sup>57</sup> Deonna and Teroni (2012b, p. 70) briefly discuss this point.

<sup>58</sup> See Comesaña (2006), Beebe (2004), and Lyons (2019) respectively.

Second, if the objector is unsatisfied by the prospect of ER borrowing a general solution to the reliabilist's generality problem, note that I will put forward an argument which directly bears on the question of how we ought to individuate our reliability assessments in Chapter 6 (§6.5). On my view, the belief-forming processes (or, more specifically, emotional competences – but more on that later) are individuated according to the subject's learned capacity to emotionally identify an object as instantiating a given evaluative property. That is, one's evaluative belief is immediately and defeasibly justified on the basis of their emotional experience if and only if that emotional experience arises from the subject's learned capacity to identify the relevant object as instantiating the relevant evaluative property. Crucially, these learned emotional capacities are fine-grained, i.e. a subject comes to emotionally learn how to identify aggressive snakes as fearsome, or spoiled foods as disgusting, or sexist comments as offensive, and so forth. Briefly, one reason for endorsing this view is that individuating *instead* on a coarse-grained basis, i.e. assessing reliability across the full range of one's emotion-types and experiences, risks objectionably ruling out cases in which an individual has an epistemically valuable emotional experience but, due to their *general* emotional unreliability, doesn't count as being sufficiently reliable in the right way so as to award *particular* emotional experiences justificatory power.<sup>59</sup>

One of the upsides of endorsing such a view is that it provides us with principled grounds for explaining why we ought to narrowly individuate emotional belief-forming processes, such that we can make sense of the verdict that the interviewer's belief that the candidate is duplicitous is *not* justified when it is formed via the emotional experience of suspicion based on the illicit bias, but it *is* (or can be) justified when it is formed via the emotional experience of suspicion based on legitimate perception of duplicitous-making features of the candidate,

---

<sup>59</sup> Later on in Chapter 6 (§6.4 and §6.5.2), I present two cases which demonstrate this point. As a brief illustration, the first of which concerns a canine behaviourist who reliably experiences fear in response to fearsome dogs but fails to reliably experience any other emotion as appropriately responsive to the presence of an evaluative property. The second case concerns a member of an oppressed social group who reliably experiences, say, anger in response to sexually offensive comments, but fails to reliably experience other emotions as appropriately responsive to the presence of evaluative properties. These will be spelled out in further detail later, but the crucial point for our purposes here is that a broad-scope ER (i.e. one which assesses for reliability across *all* of one's emotion-types and emotional experiences) would be unable to bestow the canine behaviourist's fear experiences with justificatory power, nor would be it able to allow the oppressed individual's specific anger experiences justificatory ability in virtue of their *overall* emotional unreliability. As will become clear later on, I think that this is a problematic result.

i.e. the emotional experience arises from the subject's relevant emotional competence to emotionally identify individuals as duplicitous.

Of course, the details of this proposal will be spelled out in more detail later on in the dissertation. For now, I take the availability of both various general reliabilist solutions to the generality problem and my forthcoming argument for the fine-grained individuation of processes to be sufficient to diffuse the immediate objection, such that we can explore the prospects of ER untroubled by generality worries.

### 4.3.3. A Necessity Problem

A third objection that one might level against ER concerns the putative necessity of emotional reliability for epistemic justification of evaluative belief. Specifically, one might argue that reliability of the relevant belief-forming process *cannot* be necessary for the relevant beliefs to enjoy positive justificatory status.<sup>60</sup> Let us first consider how this problem is typically presented with respect to general process reliabilism, namely, in terms of the New Evil Demon Problem (NEDP):

Imagine that, unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist. It would follow on reliabilist views that under such conditions the beliefs generated by those processes would not be justified. This result is unacceptable. The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and, therefore, that we would be just as well justified in believing what we do if the demon hypothesis were true as if it were false. (Lehrer & Cohen 1983, p. 192)

The worry is that a theory of epistemic justification which requires reliability of the relevant belief-forming process will generate the wrong result here. That is, it seems drastically counterintuitive (and perhaps even unfair), so the thought goes, to deny our demon-world

---

<sup>60</sup> On the other side of the coin, one might worry that reliability cannot be *sufficient* for epistemic justification, such that ER must be false. This objection, I contend, is significantly more challenging for the defender of ER to answer. As such, Chapter 5 (§5.2) will dedicate significant attention to this problem.

analogues with epistemic justification for their beliefs merely because the relevant belief-forming processes are unreliable in their world, and the reason for this is that the demon-world belief-forming processes are (putatively) identical in all of the relevant ways to the belief-forming processes in the normal world, which, given their reliability, *do* confer justification to the relevant beliefs. If NEDP is convincing, then reliability of the relevant belief-forming processes cannot be necessary for epistemic justification. Moreover, if this reliability is not necessary for epistemic justification, then general process reliabilism must be false. Finally, if general process reliabilism is undermined by NEDP, then perhaps, the objector may stress, we have good reason to think that ER will be similarly defeated.

I take it that there are two reasons as to why we should resist this conclusion. The first of which is that it's not obvious that the NEDP is a knock-down argument for the falsity of general reliabilism. Indeed, many advocates of reliabilism do not take NEDP to establish that reliability cannot be necessary for epistemic justification. Take Goldman's (1988) response to the challenge, for instance. According to Goldman, the reliabilist can distinguish between notions of *strong* and *weak* justification in order to dispel the result that demon-world subjects are not justified in their beliefs. For Goldman, the demon-world subjects can be said to enjoy weak justification for their beliefs in light of the fact that all weak justification requires is that the belief is formed non-culpably or blamelessly (p. 53). Strong justification, on the other hand, requires that the belief was *actually* well-formed by suitable processes. Plausibly, the demon-subjects are blameless in their beliefs (given the intuition that it would be unfair to deny them justification for beliefs formed by processes relevantly identical to normal-world subjects), and so enjoy weak justification, even if they cannot enjoy strong justification in light of the demon's intervention. Alternatively, if distinguishing between different varieties of epistemic justification strikes the reader as implausible, Baysan (2017) argues that the reliabilist need not alter their position in light of the NEDP, and suggests that appeal to a dispositionalist framework of properties can provide reliabilists with the resources to explain why the demon-world subjects *can*, in fact, be justified in their beliefs. According to Baysan, what we can say about the demon-world subjects is that their belief-forming processes plausibly instantiate the dispositional property of being reliable, but that the demon's intervention prevents this disposition from being manifested. Hence, reliabilists need not deny that the demon-world subjects are justified in their beliefs, and the counterintuitive result putatively delivered by the NEDP is dispelled.

There are, of course, many more strategies available to the reliabilist. Lyons (2013), for example, argues that the reliabilist in fact *ought* to deny positive justificatory status to some of the demon-world subjects' beliefs. Comesaña (2002), on the other hand, advances a nuanced Indexical Reliabilism which purports to diffuse the force of the NEDP, while Moon (2012b) argues that the more plausible variants of internalism *also* fall foul of the NEDP. A full discussion of these argumentative options would take us too far afield, given that the purpose of this chapter is to assess *emotional* reliabilism, but I take it that the availability of such strategies might go some way in diminishing the putative intuitive force of the NEDP for reliabilists.

A second reason to resist this challenge is that *even if* general reliabilism is undermined by NEDP, we have good reason to suppose that ER is not similarly vulnerable to this challenge.<sup>61</sup> To illustrate, consider S, a subject in the actual world, and S\*, S's demon-world counterpart. Let us suppose that S and S\* undergo psychologically identical processes of apprehending an approaching snake and its features, e.g. erratic movements, aggressive hissing, sharp fangs, and so forth. Moreover, suppose that both S and S\* have an emotional experience of fear in response to the snake. Correspondingly, both S and S\* form the evaluative belief that the snake is fearsome on the basis of their emotional experience. Now, following the example of the NEDP, we can stipulate that ER will bestow S's evaluative belief with justificatory power given S's reliability with respect to forming evaluative beliefs about fearsomeness on the basis of their fear experiences, whereas ER denies the same justificatory power to S\*'s evaluative belief given S\*'s drastic unreliability in this regard. Is this a problematic result? Plausibly not, and I take it that this is down the nature of the relationship between emotional experience and evaluative properties. Let me explain.

If we suppose that the normal-world subject and the demon-world subject are psychologically identical, then, presumably, this entails that both subjects have the same experience of the relevant non-evaluative properties of the object, i.e. both subjects experience the cognitive base of the emotional experience (e.g. the perceptual experience of the snake and its features) in the same way. Moreover, recall from earlier chapters that one plausible theory concerning the nature of evaluative properties is that they merit or make fitting particular emotional

---

<sup>61</sup> Indeed, recall that I'm taking my assessment of ER's plausibility to be logically independent from the overall truth or falsity of general reliabilism, such that, even if general reliabilism is rendered untenable by NEDP, this does not automatically mean that ER is rendered untenable accordingly.

responses. An object instantiates the evaluative property of ‘fearsomeness’, then, insofar as that object merits fear responses. Putting all of this together, then, if S is reliable with respect to forming true fearsomeness-beliefs on the basis of their fear experiences, then S reliably tracks fearsomeness-making non-evaluative properties instantiated by particular objects via perceptual experience, imaginative experience, etc. Importantly, S\* is psychologically identical to S, so S\* must also reliably track fearsomeness-making non-evaluative properties instantiated by particular objects. However, by stipulation of the NEDP example, S\* is unreliable with respect to forming fearsomeness-beliefs on the basis of their fear experiences. But, if S\* is psychologically identical to S, and thereby reliably tracks fearsomeness-making features of objects, then how *could* S\* be unreliable with respect to forming true fearsomeness-based beliefs? The only plausible answer to this question (assuming neo-sentimentalism)<sup>62</sup> is that the demon’s intervention *changes* what merits fearsomeness in the demon-world, such that what instantiates the evaluative property ‘fearsomeness’ is different in the demon-world than in the normal-world. Therefore, demon-world subjects cannot be identical to normal-world subjects in all relevant respects because the evaluative properties that they’re tracking must be different by stipulation of the case. The fact that ER does not bestow justificatory power to the demon-world subject, then, is not problematic.

So, if NEDP cannot be plausibly applied to ER, then I take it that this provides us with good reason not to abandon the necessity of reliability for emotional justification. As a further reason, recall that appeal to reliable success was what allowed ER to generate the correct result in the over-generalisation cases that plagued dogmatist versions of the justificatory thesis. The fact that building in a reliability condition into the justificatory thesis of emotion seems to significantly improve its prospects with respect to these cases (alongside the fact that ER doesn’t seem to be undermined by an analogous presentation of the NEDP) certainly seems suggestive of the idea that the reliability of forming evaluative beliefs on the basis of emotional experience is an important, and perhaps even necessary, condition for the emotions to possess immediate and defeasible justificatory power. In light of these arguments, then, I take it that the defender of ER can escape this necessity objection unscathed.

---

<sup>62</sup> Although I assume neo-sentimentalism here, note that the point of this example is compatible with other theories which draw up a substantive connection between the nature of value and emotional experience.

### 4.3 Objection: The Unreliability Problem

Now, having clarified and defended ER against some initial objections, let us move onto a further worry facing the view. One of the most straightforward means of rejecting ER is to deny the consequent of the biconditional, i.e. to deny that forming evaluative beliefs on the basis of emotional experience is a reliable process of belief-formation. So, if it's true that reliability is necessary for the emotions' ability to immediately and defeasibly justify evaluative belief (as argued for in §4.2.3), and if forming evaluative beliefs on the basis of one's emotional experience is *in fact* an unreliable belief-forming process, then evaluative beliefs formed on the basis of one's emotional experiences cannot be justified. Call this objection *the unreliability problem*.

The idea that emotional experience is unable to meet a sufficient threshold of reliability in order to perform a positive epistemic role is a familiar one. Indeed, recall from Chapter 1 (§1.1) that opposing the justificatory thesis of emotion is the idea that emotional influence makes us epistemically *worse off* than we would otherwise be; that emotions bias and cloud our perception as to the way things are. It certainly seems to be a common folk intuition that hot-headed emotional reasoning is likely to lead us astray, while cool and cerebral rationality is the reliable pathway to truth. Moreover, we can look to the philosophical literature for comments suggestive of this intuition, such as Tappolet's (2016) claim that "it is far from clear that the relation between emotions and concerns is a reliable one. More often than not, emotions misfire..." (p. 42). An immediate challenge for ER, then, is to defend the reliability of forming evaluating beliefs on the basis of emotional experience against this appearance of unreliability.

To be clear, the unreliability problem cannot amount to the simple claim that emotions *sometimes* misfire, for advocates of ER will be perfectly happy to accept this possibility given the caveat that emotional experience only defeasibly justifies evaluative belief on this model. In order to do any substantive damage to ER, then, the unreliability problem must be presented not just as the worry that emotions misfire in unideal conditions, but instead as the worry that emotional experience is significantly or systematically unreliable, such that it is uniformly incapable of conferring any degree of epistemic justification. It is not the problem that there are token instances in which emotional experiences fails to generate true evaluative



beliefs, but rather, it is the problem that emotions *so often* misfire such that they do not meet the reliability precondition required for ER to obtain.

In what follows, the question I will endeavour to answer is whether there exists any plausible theoretical support for the idea that are emotions are systematically unreliable. I take there to be three possible arguments one might present in favour of such unreliability: (1) emotional experience is highly cognitively penetrable, and this high rate of cognitive penetration is the source of the emotions' unreliability, (2) the complexity and obscurity of value concepts severely hinders the emotions' ability to accurately track evaluative information, and (3) the possibility of reliable emotional dispositions in virtuous agents is ruled out by: (i) situationist challenges, and/or (ii) the absence of any compelling feedback system with respect to emotion and evaluative belief. In what remains of this chapter, I will elucidate and ultimately reject each of these arguments in turn.<sup>63</sup>

#### 4.3.1 The Cognitive Penetrability of Emotion

One possible explanation of the emotions' widespread unreliability might be that the content of our emotional experiences appears to be significantly susceptible to being unduly influenced by background cognitive states, such as beliefs or desires. For example, a subject's unjustified belief that members of a particular ethnic group are untrustworthy will plausibly cause them to inappropriately experience fear or suspicion towards those individuals; a subject's desire to rush home in time for the series finale of their favourite show will plausibly cause them to inappropriately experience anger towards leisurely pedestrians blocking their path. To put it more formally, one might think that emotional experience is *cognitively penetrable* to a significant degree.<sup>64</sup>

---

<sup>63</sup> At this point, one might object that the putative unreliability of emotion is a matter of brute empirical fact, such that, *even if* we cannot find a plausible avenue of theoretical support for the systematic unreliability of emotion, the fact of the matter remains that forming evaluative beliefs on the basis of one's emotional experience is *just* an unreliable process of belief-formation. Let me clarify two points in response. First, the burden of proof is on the advocate of the unreliability problem to provide widespread empirical support for the idea that emotional experience is unreliable *to this extent*, i.e. that emotions are insufficiently reliable across all agents and contexts. Second, in the absence of this support, the *further* absence of a convincing theoretical explanation for this idea must count against the plausibility of the unreliability problem.

<sup>64</sup> Interestingly, we've already discussed putative cases of cognitive penetration. Recall the focus on Markie's gold prospector in the previous two chapters. Plausibly, what's going on in the case of Markie's wishful prospector is that their desire to discover gold *cognitively penetrates* their visual

In the interest of clarity, let us formally define cognitive penetration in terms of the following from Vance (2014):

A visual experience is cognitively penetrable with respect to some content or character *c* if and only if two subjects (or the same subject at different times) can differ with respect to whether their experience has *c*, and the difference is the result of a causal process tracing back to a non-visual psychological state of the subject, where we hold fixed between the two subjects (or one subject at different times) the following: (i) the stimuli impacting their sensory receptors, (ii) the subject's spatial attention, and (iii) the conditions of the subjects' sensory organs. (p. 259)

On this definition, an experience is said to be cognitively penetrable if the content of that experience is liable to being non-trivially influenced by the subject's background cognitive states.<sup>65</sup> For an illustration of a putatively cognitively penetrated visual experience, consider the following case from Siegel (2012):

Jill believes, without justification, that Jack is angry with her. The epistemically appropriate attitude for Jill to take toward the proposition that Jack is angry at her is suspension of belief. But her attitude is epistemically inappropriate. When she sees Jack, her belief makes him look angry to her. If she didn't believe this, her experience wouldn't represent him as angry. (p. 209)

Siegel's case appears to fulfil the conditions for cognitive penetrability outlined in the definition above. Without the belief that Jack is angry with her, and holding fixed all other relevant conditions, Jill would not visually experience Jack's facial expression as angry. Thus, it seems to be the case that Jill's background cognitive state, her belief that Jack is angry with her, is penetrating her visual experience.<sup>66</sup>

---

experience, i.e. the visual experience with the contents 'this yellow pebble is gold' can be causally and non-trivially traced back to the prospector's strong desire to discover gold.

<sup>65</sup> See Macpherson (2012) and Siegel (2012) for similar definitions.

<sup>66</sup> The extent to which visual experiences are cognitively penetrated, if they can be at all, is largely an empirical matter. Some have argued that visual experiences are, in fact, impervious to penetration from cognitive states; that cases like Siegel's can be analysed and re-described in ways that do not require the possibility of cognitive penetration. I will assume here that cognitive penetration of visual experiences is empirically possible, for my interest lies in the epistemic implications of cognitive

What makes cognitive penetration relevant to this discussion is the epistemic challenge it raises for accounts of justification. Some, including Siegel, have suggested that this putative cognitive penetrability of visual experiences may have negative epistemological ramifications in certain cases; that such a feature might threaten or at least limit the justificatory ability of certain visual experiences. In an ordinary case, i.e. a case in which Jill's visual experience is *not* cognitively penetrated, it may be plausible to suggest that Jill's visual experience of Jack's facial expression defeasibly justifies her belief that he is angry with her. However, because the content of Jill's visual experience has been influenced by her antecedent unjustified belief that Jack is angry, one might worry that this influence effectively bars that visual experience from doing any justificatory work. That is, crucially, cognitive penetration does not just serve as a defeater for the *prima facie* justification ordinarily conferred by visual experience, but rather it blocks the experience from performing *any* justificatory role.<sup>67</sup>

So, if we have reason to think that the justificatory power of visual experiences may be threatened or limited by their putative cognitive penetrability, then we may have reason to think that emotional experiences are similarly liable to this possibility. On the question of whether the notion of cognitive penetrability can be applied to emotional experience, Vance (2014) has presented the following putative definition as a direct analogue to his initial criteria for the cognitive penetration of visual experiences:

Emotional states are cognitively penetrable with respect to some content or aspect of phenomenal character *c* if and only if two subjects (or the same subject at different times) can differ with respect to whether their emotional states have *c* as the result of a causal process tracing back to a cognitive state of the subject, holding fixed between the two subjects (i) the stimuli impacting their sensory receptors, (ii) the subjects' spatial attention, and (iii) the conditions of the subjects' sensory and emotional organs. (p. 270)

---

penetrability. Note, moreover, that a variety of empirical studies have been conducted which suggest that cognitive penetrability of this sort we're interested in here is possible. See, for example, Levin and Banaji (2006) and Perky (1910). See also Macpherson (2012) for a helpful survey and analysis of this evidence.

<sup>67</sup> Chudnoff (2020) refers to this as the 'Downgrade Thesis' (p. 2), i.e. the view that cognitively penetrated, or "hijacked", experiences are incapable of lending immediate and defeasible justification to belief because they fall below a certain epistemic baseline.

Reconsider the emotional examples cited above. In both cases, the content of the emotional experience would be different in the absence of the influencing cognitive state. Upon encountering the individual from a particular ethnic group, the xenophobe's emotional experience would not be one of suspicion if they didn't already have the unjustified antecedent belief that members of this group were untrustworthy; one's emotional experience would not be one of anger towards the leisurely pedestrians if they didn't already have the antecedent desire to get home quickly. Indeed, on this definition, it seems like a great many of our emotional experiences are going to turn out to be cognitively penetrated. Holding all the factors fixed in the above definition, the content of my emotional experience will often differ from yours due to the fact that you and I have different background desires, beliefs, etc. I may experience contempt in response to something that you find amusing, you may experience sadness in response to something that brings me great joy, and so forth. Thus, it might be suggested that this high cognitive penetrability of emotional experience is a precisification of the intuition that underlies the unreliability problem. Emotions are so often misleading and unreliable because the content of emotional experience is heavily vulnerable to infection from a host of psychological pollutants.

Now, an immediate thing to say about this suggestion is that the mere fact that emotional experience bears a higher rate of cognitive penetration is insufficient to arrive at the conclusion that the emotions are systematically unreliable. A further step is required, and that is to show that *most* instances of emotional cognitive penetration are epistemically damaging. That is, it would have to be argued that the surge in influence from our background beliefs and desires makes us *significantly* epistemically worse off than we would be otherwise in order to decrease the veridicality ratio of emotional experience to a level insufficient for reliability. This seems to be implausibly strong for two reasons.

Firstly, very few in the perception literature would be comfortable endorsing the claim that cognitive penetration of visual experience (if and when it occurs) always or for the most part suppresses the potential justificatory power of the experience. Discussion of the epistemic implications of cognitive penetration tends to be more nuanced than indiscriminate claims like this. Siegel (2012), for instance, argues that the justificatory power of visual experience is extinguished by cognitive penetration *only* when an epistemically problematic circularity is introduced (p. 209). Jill's visual experience of Jack's angry facial expression fails to justify

her belief that Jack is angry with her because the causal root of her having that visual experience *is* her unjustified belief that Jack is angry with her. Jill, so the argument goes, effectively enters into a feedback loop with respect to her justification. To generalise, what is epistemically illicit about (some) cases of cognitive penetration is that the causal origin of the process is the unjustified belief that X is F. This belief then causes the subject to visually experience X as F. The subject then bases her belief that X is F on her visual experience that X is F. Because the belief that X is F is thereby plausibly understood as being *based on itself*, Siegel takes cognitive penetrability to extinguish the justificatory power of visual experiences only when it leads us into these epistemically vicious feedback loops.

Alternatively, Lyons (2011) proposes that the root of the epistemic challenge of cognitive penetration concerns not epistemically vicious circularity but the question of *reliability*. According to Lyons, cognitive penetration inhibits the epistemic function of visual experience *only* when the penetrating state curtails the overall reliability of the perceptual system. On this account, what is epistemically worrisome about cases like Siegel's is that the cognitive penetration of Jill's visual experience – the manipulation of her visual experience by her antecedent unjustified belief – biases her perception such that she would visually experience Jack as angry *regardless* as to whether this were true (p. 301). Certain instances of cognitive penetration render us unresponsive and insensitive to the way things are, and *this* is what is epistemically problematic. The insensitivity to facts caused by certain instances of cognitive penetration decreases the likelihood of our beliefs turning out to be true, and this threat of unreliability is what inhibits the justificatory power of visual perception.

Crucially, in both Siegel's and Lyons' account, the epistemic outcome of cognitive penetration is judged locally, rather than globally across all cases. Therefore, because there is very little in the way of evidence for the claim that cognitive penetration of visual experience uniformly or for the most part inhibits the justificatory power of the experience, it seems that we have good reason to believe there is analogously little evidence for the claim that cognitive penetration of emotional experience is chiefly epistemically vicious.

Secondly, considering emotional cognitive penetration directly, it seems implausible to claim that cognitive states consistently corrupt my emotions for the worse. For example, my background beliefs about the technical expertise involved in performing certain pieces of music may penetrate my emotional system such that I reliably experience appropriate

admiration in response to admirable performances of those pieces; my desire to preserve my bodily integrity may penetrate my emotional system such that I more reliably experience fear in response to physically threatening objects. In cases like these, and plausibly many more, the cognitive penetrability of emotional experience is not epistemically damaging, but instead appears to be epistemically beneficial; it places me in a better epistemic position with respect to grasping my evaluative surroundings than I would be in without the cognitive influence.

Therefore, because the cognitive penetrability argument can only succeed in establishing the emotions' systematic unreliability by endorsing the claim that most instances of emotional cognitive penetration are epistemically vicious, and because this claim seems to be implausibly strong, I conclude that this line of argument fails to support the unreliability problem, and thereby fails to undermine ER.

### 4.3.2 The Obscurity of Value

A second possible argument in favour of the systematic unreliability of emotional experience concerns not the nature of the emotional experience itself, but the nature and obscurity of evaluative information and properties. To contrast, return to the case of visual experience. Our visual systems seek to gather non-evaluative information about our surroundings, tracking properties of colour, shape, motion, etc. These visual properties are, at least in normal cases, straightforwardly grasped by perceivers. On the other hand, evaluative information which constitutes the core relational themes of anger, envy, pride, shame, indignation, etc. is relatively more obscure, and consequently more difficult for ordinary perceivers to grasp.

In elucidating this claim, one might point to one of the following two features of our experience with value concepts as putative indicators of the comparative obscurity of value. First, one might identify the widespread disagreement pervading our value discourse as such an indicator. It seems to be the case that everyday experience lends support to the claim that we often encounter disagreement in our value judgments.<sup>68</sup> For instance, I might judge a

---

<sup>68</sup> A stronger claim might be that evaluative concepts do not just *often* incite substantive disagreement (that is, disagreement concerning the content of the concept), but rather that they *essentially* admit such disagreement. That is, one might claim that evaluative concepts are *essentially contestable*, and that this feature indicates their obscurity. See Gallie (1955).

particular joke to be offensive, whereas somebody else might judge it to be amusing; I might judge somebody's action to be admirable, whereas somebody else might judge it to be dishonourable, and so forth. Disagreements like these happen often, or at least significantly more so than disagreement concerning, say, the properties of colour and shape tracked by our visual systems. Under ordinary conditions, normal observers do not tend to encounter the same (or even a similarly high level of) disagreement concerning whether the colour of the door is blue or red, or whether the shape of the water bottle is round or square. One might take this comparatively widespread disagreement to be suggestive of the fact that there is little in the way of common understanding of the evaluative information which constitutes our concepts of value.<sup>69</sup> In turn, we might then take this absence of any robust consensus as indicative of the obscurity of value.

Relatedly, one might identify the lack of any determinate boundaries between value concepts as a potential indicator of their obscurity. Again, to illustrate against the background of perceptual experience, consider properties tracked by our visual systems. Typically, it is not the case that ordinary observers suffer any substantial confusion about the boundaries between what is green and what is red; what is round and what is square; what is moving and what is stationary. However, this does not seem to be true with respect to evaluative properties. It is significantly less clear what the demarcating boundaries are between evaluative concepts such as 'offensive' and 'insulting', 'shameful' and 'embarrassing', 'contemptible' and 'disgusting', and so forth. Thus, our inability to accurately identify determinate boundaries between various value concepts might be understood as suggestive of the obscurity of value.<sup>70</sup> Therefore, because the evaluative information putatively sought after

---

<sup>69</sup> As a point of qualification: it seems at least *prima facie* true that different evaluative concepts bring about varying levels of disagreement. Perhaps it may be difficult for a diverse pool of observers to agree unanimously on what the core relational theme of 'shamefulness' consists in, but it seems significantly more likely that that pool of observers will agree on, say, the wrongness of particular actions such as murder or incestual relationships. This discrepancy of disagreement across evaluative concepts might suggest that *some* evaluative information is obscurer than others, such that *some* of our emotion-types are correspondingly less likely to be reliably responsive to the presence of evaluative properties than others. This thought might prompt us to consider adopting a finer-grained approach to making assessments of reliability (as I suggested in §4.2.2), but I set this point aside for now, given that my purpose here is to argue that the obscurity of evaluative information is not a plausible explanation of the systematic unreliability of emotional experience.

<sup>70</sup> Similar points concerning the indeterminacy of certain evaluative concepts have been made by others in the literature. Szigeti (2015), for example, discusses the indeterminacy of the core relational theme of guilt (p. 15), while D'Arms and Jacobsen (2003) point out that there is "considerable internal disagreement within the tradition: for instance, over how to differentiate guilt from remorse, and regret from... agent-regret" (p. 134).

by emotional experience is so relatively obscure, the advocate of the unreliability problem might take this to be a plausible diagnosis and theoretical backing for the systematic unreliability of emotional experience. The opacity of evaluative information hinders our emotions' ability to accurately track value, and, therefore, forming evaluative beliefs on the basis of one's emotional experiences cannot constitute a reliable belief-forming process.

The problem with this argument is that it involves opening the door to a widely encompassing and pernicious scepticism with respect to our access to evaluative information. After all, if it is the case that evaluative information is so obscure such that it severely hinders our emotions' ability to accurately track value, then, plausibly, the obscurity of value *also* theoretically disables all other means we have of accessing such information. Our inability to grasp value seriously calls into question the accuracy of both our reasoning about value and the body of our evaluative beliefs and judgments. If the advocate of the unreliability problem wishes to push this line of argument in support of their objection, then it seems that they must allow this pervasive scepticism to creep into their account; they must concede that the obscurity of value blocks any secure epistemic access we have to the evaluative realm. This is a heavy theoretical burden to bear. No secure access to the evaluative reality threatens the existence of all sorts of evaluative knowledge and understanding we intuitively take ourselves to have, including moral and aesthetic knowledge.

In summary, then, because diagnosing the systematic unreliability of emotional experience by appeal to the opacity and obscurity of evaluative information involves opening the door to serious scepticism about evaluative knowledge, and because of the significant theoretical costs involved in doing so, I argue that this line of argument fails to support the unreliability problem, and thereby fails to undermine ER.

### **4.3.3 Against the Reliable Emotional Dispositions of Virtuous Agents**

Now, in defence of ER, one might argue that there *are* circumstances in which forming evaluative belief on the basis of emotional experience is a reliable process of belief formation, namely, circumstances concerning the emotional dispositions of virtuous agents. The possession of the virtues, i.e. particular moral and intellectual character traits, plausibly involves certain dispositions to, as Pelsler (2011) puts it, "have the right emotions towards the right objects at the right time and to the right degree" (p. 14). The virtuous agent, so the



thought goes, will feel the right degree of anger in the right circumstances involving an illegitimate wrong; she will feel the right degree of sadness in the right circumstances involving loss, and so forth. In reliably experiencing her emotions in the appropriate circumstances, the virtuous agent may plausibly be understood as possessing emotional dispositions which accurately track value. Therefore, if it is true that *some* people possess at least *some* virtues,<sup>71</sup> then there are cases in which forming evaluative belief on the basis of emotional experience is a reliable process of belief formation.<sup>72</sup> Indeed, it is when the emotions are functioning properly in the virtuous agent that they fulfil their epistemic role and vindicate ER.

A third potential way in which the advocate of the unreliability problem might attempt to source theoretical support for their position concerns eliminating and disavowing the possibility of reliable emotional dispositions in virtuous agents. That is, one might attempt to dismiss this argument, and consequently reinforce the systematic unreliability of emotion, in one of two ways: (i) appeal to situationist critiques of virtue ethics in order to challenge the possession of virtue, and (ii) appeal to the absence of any robust and compelling feedback system with respect to evaluative belief in order to challenge the idea that virtues can be developed in the first place. This section will reject both of these arguments in turn.

First of all, the advocate of the unreliability problem might note that there is an affluence of empirical evidence which suggests that possession of the virtues is incredibly infrequent, if attainable at all. Situationist psychological studies have called the existence of virtues into question by showing that merely manipulating the non-moral, situational factors of a scenario has a significant influence on the moral behaviour of others.<sup>73</sup> These studies, roughly, work by slightly altering a non-moral variable within a situation and find that the change in this variable reliably produces a change in the moral behaviour of subjects. One example of this is Isen and Levin's (1972) study in which ordinary pedestrians walk into a payphone. Some of these subjects walk in to find a dime, and others do not. Regardless as to whether a dime was

---

<sup>71</sup> Most virtue ethicists, in opposition to the traditional Aristotelian view, allow for the possibility of an agent's possessing some virtues without possessing others; that one can possess the virtue of generosity without possessing the virtue of honesty, for example. However, see Badhwar (1996) for a nuanced discussion of how the virtues systematically relate to one another.

<sup>72</sup> In opposition to a view of this sort, however, one might argue that the virtuous agent simply will not form evaluative beliefs on the mere basis of their emotional experiences. See Brady (2010) for a view of this sort.

<sup>73</sup> See Harman (1999) and Doris (1998)

found, when the subjects left the payphone, they came across a pedestrian who dropped a pile of papers in the street. Researchers monitored which of the subjects stopped to help the distressed passer-by and found that those who found the dime in the payphone were significantly more likely to help.<sup>74</sup> What these challenges show is that manipulating situational factors appears to alter the moral behaviour of others, such that these situational factors have a greater explanatory power with respect to our moral actions than the putative existence of moral virtues. Therefore, in light of these challenges, it is thought that the existence of robust, persisting character traits ought to be called into question.

Now, recall from Chapter 1 (§1.1) that one important motivation for endorsing the justificatory thesis of emotion was that it promised a substantive epistemic yield, i.e. it promised a simple and naturalised explanation as to how we obtain justified evaluative belief, such that we could annul sceptical worries pertaining to the acquisition of justified evaluative belief. In order for the appeal to virtuous agency to fully vindicate ER, it would have to be shown that the possession of virtue is frequent enough such that the reliabilist conception of the justificatory thesis of emotion could retain this advantage, i.e. it would have to be shown that possession of virtue is frequent enough such that, for many of us, forming evaluative belief on the basis of emotional experience is a reliable process of belief formation, such that many of us, rather than just the virtuous minority, can enjoy epistemic justification for our evaluative beliefs on the basis of our emotional experiences. The advocate of the unreliability problem, in amassing the conclusions drawn from these situationist studies as evidence for absence of regular, persisting virtues, could argue that appeal to the reliable emotional dispositions of virtuous agents is too insubstantial to do any philosophical work in this regard. Possession of virtue is much too infrequent in order to vindicate ER in a way that retains the original attractions of the justificatory thesis of emotion.<sup>75</sup>

However, this line of argument – disavowing the possibility of virtuous agency on account of situationist challenges – suffers a similar dialectical problem to the one facing both the cognitive penetrability argument and the obscurity of value argument. This argument hinges entirely on the philosophical power of situationist challenges and, thereby, on the falsity of

---

<sup>74</sup> For further situationist experiments along these lines, see Milgram (1974) and Darley and Batson (1973).

<sup>75</sup> Carter (2020) appears to have a worry of this sort with respect to a reliabilist conception of the justificatory thesis of emotion, or, what he refers to as “Epistemic Perceptualism” (p. 1237).

popular strands of virtue ethics. Again, this seems like heavy theoretical baggage to carry. Virtue ethicists certainly do not take the conclusions of situationist studies to dismantle their position. Sreenivasen (2002), for example, has argued that situationist challenges fail to undermine virtue ethics due to the fact that there is a fundamental mismatch between the virtue ethicist's conception of a character trait, and the particular conception of a character trait targeted and putatively threatened by situationist challenges. According to Sreenivasen, situationist challenges rest on a deficient "operationalised conception of a 'character trait'" (p. 47), and this inadequacy is rooted in their failure to take into account relevant factors of the agent's situation, such as whether they themselves construe their situation as pertinent to that character trait, how relevant their behaviour actually is to the character trait in question, etc. On a similar note, Annas (2005) has argued that situationist challenges fail due to the fact that these challenges misguidedly target versions of virtue theories which identify character traits as mechanical, unthinking reflexes. According to Annas, most in the virtue ethics tradition instead conceive of character traits as (or involving) cognitive capacities to act on the basis of reasons (p. 637). In the payphone study, the fact the agents opted not to help the struggling confederate is insufficient to decisively show that those agents lacked the virtue of, say, generosity. Virtues need not manifest themselves as habitual, automatic responses to certain moral triggers. As a third alternative, Sabini and Silver (2005) argue directly against the payphone situationist challenge, contending that picking up papers simply isn't a significant enough moral action to provide us with any substantive information about an agent's character. In this case, according to Sabini and Silver, an agent's being in a bad mood (or, to be precise, an agent's not being in a good mood relative to those who found the dime) may be a perfectly legitimate reason for them not to pick up the papers.

It is not within the scope of this chapter to review and evaluate all possible means of responding to these challenges. Rather, the point here is to show that simple appeal to situationist experiments – without due consideration of possible responses made on the behalf of virtue ethicists – is insufficient to decisively rule out the possibility of possession of virtue, and consequently, reliable emotional dispositions of virtuous agents. If the advocate of the unreliability problem wishes to push this line of argument, they would have to systematically trawl through each of these responses and show that none of them can be successful. While not impossible, this places a significant and problematic explanatory burden on the advocate of the unreliability problem. Thus, brute appeal to situationist

experiments in order to disavow the possibility of reliable emotional dispositions in virtuous agents fails to support the unreliability problem, and thereby fails to undermine ER.

Finally, the advocate of the unreliability problem might attempt to disavow the possibility of reliable emotional dispositions in virtuous agents by challenging the idea that virtuous agents *can* develop reliable emotional dispositions in the first place. According to most Aristotelian strains of virtue ethics, virtues are not simply embedded into one's character from birth, but instead are *developed* over one's life through conscious habituation and practice. If it is the case that the possession of virtue involves the possession of the certain emotional dispositions, then these emotional dispositions are also plausibly understood as being cultivated and refined through the relevant forms of practice and training. However, one might argue that this notion of practicing virtue is implausible due to the fact that there is no appropriate feedback system in place to correct our emotional dispositions and thus facilitate such training.

To illustrate this point, reconsider sense perceptual experience. As we go about our lives, it seems that we regularly receive 'feedback' on how accurately our perceptual experience is representing the world around us. If I'm driving a car, and my perceptual system is misconstruing properties of shape and motion to me, I'll straightforwardly receive feedback on this failure in the form of my crashing into obstacles; feedback which then prompts me to recalibrate my perceptual system to improve its accuracy and reliability. It's less clear what analogous sense of feedback we might get with respect to emotional experience and evaluative information. The closest thing we seem to receive in terms of feedback in the evaluative case is disagreement with others. Suppose I experience amusement towards a certain joke, and thereby judge it to be amusing, whereas you fail to do so. Unlike feedback received in the perceptual case, it's not at all clear that this mere disagreement will compel or even prompt me to recalibrate my emotional system in order to bolster its accuracy. I can explain away your lack of amusement (perhaps by appeal to the background beliefs and desires mentioned in the cognitive penetration argument) without admitting the possibility of my evaluative mistake, such that mere fact of our disagreement does not compel me to reassess my evaluative perception of amusement properties in the same way that, e.g.

crashing my car compels me to reassess my visual perception of properties of motion and space.<sup>76</sup>

In summary, then, the absence of any similarly compelling feedback system with respect to evaluative properties means that I cannot check my evaluative perception against the evaluative reality, whereas it seems that I can rather straightforwardly check my visual perception against the non-evaluative reality. The advocate of the unreliability problem might cite this absence of feedback as a barricade for the potential of ‘training’ emotional dispositions, and thereby, as a means of explaining away the possibility of reliable emotional dispositions in virtuous agents.

For the sake of argument, let us assume that a compelling evaluative feedback system is necessary for the development of reliable emotional dispositions. It seems to me that the defender of ER can respond to this objection in one of two ways. First of all, they might point out that while *some* evaluative properties admit little in the way of compelling feedback, this is not the case for all evaluative properties. Take fearsomeness, for example. If my emotional system is failing to accurately track properties of fearsomeness, a plausible consequence of this is that I’ll fail to evade objects which constitute genuine threats to my bodily integrity, and thereby endure greater and more frequent injury than I would do were my emotional system to correctly track those properties. Similarly, if my emotional system is failing to accurately track properties of disgustingness, a plausible consequence of this is that I’ll fail to evade contaminated objects, and thereby suffer illness and infection more frequently than I would do were my emotional system to correctly track those properties. Hence, it does seem as though there is a compelling feedback system in place with respect to evaluative properties. Note, of course, that this response remains concessive to the advocate of the unreliability problem insofar as it allows that only *some* evaluative properties admit suitable feedback, and thereby that only the relevant corresponding emotional experiences can be ‘trained’ to the level required for virtuous agency.

As an alternative, then, the defender of ER might push back against the idea that all instances of disagreement are uniformly weak sources of feedback. Certainly, it seems true that I am

---

<sup>76</sup> Kauppinen (2013) makes a similar point concerning the lack of any substantive feedback system specifically with respect to moral judgments.

not *compelled* to recalibrate and re-examine my emotional dispositions every time I encounter somebody who disagrees with my evaluative judgment. However, it seems to be equally true that I regularly experience the strong inclination to revise my evaluative judgment when it comes into conflict with the evaluative judgments of those that I respect. Suppose, again, that I judge a particular joke to be amusing, whereas my mother (i.e. somebody that I consider to be virtuous) judges the joke to be cruel and offensive. Further, suppose that she confronts me about my amusement and asserts her strong disagreement of my judgment. It seems quite straightforwardly true that in *this* instance, i.e. an instance in which my interlocutor is somebody that I consider to be a role model, my being faced with this disagreement *will* prompt me to re-evaluate the fittingness of my emotional experience. Receiving evaluative feedback in the form of disagreement or disapproval from those we respect is a central part of our moral development, and thereby, seems to me to be a suitable candidate for the type of feedback system putatively required for the development of reliable emotional dispositions in virtuous agents. Tarring all forms of disagreement with the same critical brush overlooks this integral part of our moral and evaluative development, and, as such, I argue that this objection fails to support the unreliability problem, and thereby fails to undermine ER.

#### 4.4 Concluding Remarks

The purpose of this chapter has been to construct and make tenable a reliabilist account of emotional justification which identifies the source of the emotions' justificatory power in the fact that forming evaluative beliefs on the basis of emotional experience is a *reliable* process of belief formation.

§4.1 began by identifying various authors in the philosophy of emotion literature that seem to attribute epistemic importance to the emotions' reliable connection to evaluative properties. In §4.2, I then suggested that, in light of this, a plausible contender for an attractive account of emotional justification is a view based on process reliabilism. From this, I identified the target view as emotional reliabilism (ER), according to which emotional experience is capable of immediately and defeasibly justifying an evaluative belief *e* if and only if forming evaluative beliefs on the basis of emotional experience is a reliable process of belief-formation. I then attempted to motivate this view by highlighting its advantages and defending it against three preliminary objections. The first of which stated that the putative

absence of a required causal interaction between emotional experience and evaluative properties undermined the plausibility of ER (§4.2.1), the second levelled a generality problem against the view (§4.2.2), and the third objected that reliability of the relevant belief-forming process cannot be necessary for epistemic justification (§4.2.3). I argued that none of these objections were strong enough to undermine ER.

In §4.3, I then presented what I take to be the most conspicuous objection facing ER, namely, the unreliability problem. The unreliability problem argued that emotional experience *cannot* immediately and defeasibly justify evaluative belief in virtue of the fact that forming evaluative beliefs on the basis of emotional experience is *not* a reliable process of belief formation; that emotional experience is systematically unreliable such that this widespread unreliability renders it incapable of acting as a source of epistemic justification. In order to investigate the prospects of this problem, I then presented what I take to be the three most plausible arguments in favour of such systematic emotional unreliability, and argued that all three of these arguments failed. The first argument (§4.3.1) suggested that the emotions' cognitive penetrability explained their significant unreliability, i.e. the fact that the contents of our emotional experiences are significantly vulnerable to influence by our background beliefs and motivational states is what explains their deflated veridicality ratio. I argued that, even if our emotional experiences are so vulnerable, this is insufficient to get us to the conclusion of systematic unreliability, given that cognitive penetration can also improve the epistemic status of certain experiences. The second argument (§4.3.2) focused on the comparative obscurity and ineffability of evaluative information, and argued that this obscurity rendered the emotions incapable of reliably tracking evaluative properties. I suggested that this argument fails in virtue of its opening the door to a pernicious form of scepticism; that if we concede that the obscurity of value inhibits emotional access to evaluative properties, we must concede that all other forms of access to these properties fail. Finally, the third argument presented in favour of the unreliability problem (§4.3.3) attempted to disavow the possibility of reliable emotional dispositions in virtuous agents by appealing to: (i) situationist challenges, and (ii) the absence of any compelling feedback system with respect to emotional detection of evaluative properties. I argued that, first, situationist challenges are not sufficient to undermine the possibility of virtue, and second, we can make sense of such a compelling feedback system. The unreliability problem, I conclude, has no plausible theoretical support.

In what follows, however, we'll see that ER's success with respect to the specific challenges mentioned in this chapter does not guarantee overall plausibility for the view. Chapter 5 will level a serious objection against ER which, I believe, cannot be so easily overcome.



## Chapter 5

### The Problem with Emotional Reliabilism

#### 5.1 Introductory Remarks

The primary concern of the previous chapter was to present and defend emotional reliabilism against some preliminary objections. Recall that we're understanding the view as follows:

Emotional Reliabilism (ER): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if forming evaluative beliefs on the basis of emotional experience is a reliable process of belief formation.

For ER, the justificatory ability of emotional experience requires that forming evaluative beliefs on the basis of emotional experience tends to produce more true beliefs than not (and, moreover, this requires that the emotional system reliably generates true evaluative contents). Thus far, I have rebutted four objections facing ER. First, I rejected the idea that there is a problematic lack of an appropriate causal connection between emotional experiences and evaluative properties (§4.2.1). Second, I argued that the emotional reliabilist has the resources available to her to answer generality-style challenges facing reliabilist theories of justification (§4.2.2). Third, I defended ER against the objection that reliability of the relevant belief-forming process is not necessary for epistemic justification (§4.2.3). Finally, I argued against an objection which contends that forming evaluative beliefs on the basis of emotional experience is an unreliable process of belief formation given the unreliability of our emotional systems. I endeavoured to show that there is no plausible way of arguing for epistemically significant emotional unreliability (§4.3).

Now, while the view may be capable of withstanding *these* particular objections, this does not spell smooth sailing for ER. In this chapter, I will argue that there is a further problem facing the reliabilist approach to emotional justification which cannot be so easily overcome. More specifically, I argue that emotional clairvoyance-type cases establish that the reliability of forming evaluative beliefs on the basis of emotional experience cannot be sufficient for epistemic justification. From a discussion of this objection, I conclude that ER cannot provide a plausible explanation as to how emotional experience can immediately and defeasibly justify evaluative belief.

The structure of this chapter is as follows. In §5.2, I begin by setting out how clairvoyance objections are typically levelled against general reliabilist views of justification, before presenting an analogous emotional case in order to demonstrate the insufficiency of emotional reliability for justification. §5.3 then raises the possibility of modifying ER in order to evade the bite of clairvoyance-type counterexamples while also striving to preserve the core reliabilist thesis of the original view. Accordingly, I present and evaluate two alternative emotional reliabilist views inspired by modified reliabilist views provided in the literature: (i) Evidentialist Emotional Reliabilism (§5.3.1) and (ii) Agent Emotional Reliabilism (§5.3.2). Ultimately, I argue that neither of these views can provide a successful framework for the justificatory thesis of emotion. Finally, I provide concluding remarks in §5.4, and sum up the central lessons to be learned from the discussions of Chapters 4 and 5.

## 5.2 Objection: The Insufficiency of Emotional Reliability

Let me begin here by presenting and motivating a standard objection facing reliabilist accounts of justification, namely, the objection that it's not at all clear that a reliable process of belief-formation *alone* is sufficient for epistemic justification. Here's a popular way of presenting this challenge from BonJour (1980):

Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kind of subject matter. He possesses no evidence or reasons of any kind for or against the generally possibility of such a cognitive power, or for or against the thesis that he possesses it. One day, Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable. Is Norman epistemically justified in believing that the President is in New York City? (p. 62)<sup>77</sup>

---

<sup>77</sup>Another way of presenting this challenge is in terms of Lehrer's (2000) Mr. Truetemp case. Briefly, this case involves a subject, Mr. Truetemp, who is unknowingly fitted with a small machine which accurately computes the ambient temperature of whatever environment Mr. Truetemp is in. Accordingly, this device immediately causes Mr. Truetemp to form true beliefs about the temperature of his surrounding environment. Plausibly, so the thought goes, despite being maximally reliable in forming true beliefs about the ambient temperature of his surroundings, Mr. Truetemp is not justified in these beliefs, nor do these beliefs count as instances of knowledge. Given the similarities that this

According to Bonjour, a case like this meets the reliability requirement for process-reliabilist theories of epistemic justification. That is, so long as Norman is reliable in forming beliefs about the President's location on the basis of his clairvoyant experience, i.e. so long as that process tends to output more true beliefs than false ones, Norman is justified in his belief that the President is in New York. Many think that this result is implausible; that Norman *cannot* be justified in his belief that the President is in New York. In effect, then, what cases of this sort do is level a similar over-generalisation charge against process reliabilism to the kind of challenge we saw levelled against dogmatist accounts of justification. That is, cases like that of Norman's clairvoyance putatively establish that reliabilism over-generalises to cases in which we ought *not* bestow positive justificatory status to the belief in question, just like basic phenomenal conservatism putatively over-generalised to experiences like that of Markie's gold prospector. Reliability of the relevant belief-forming process alone, then, cannot be sufficient for epistemic justification.

The significance of clairvoyance cases is often taken to go beyond this negative claim, however. Many take the intuitive verdict delivered by these cases to illuminate a further necessary condition on epistemic justification. Internalists and externalists tend to spell out this condition in different ways. One explanation typically appealed to by internalists, for example, is the idea that the presence of epistemic justification seems to require the subject to have an appropriate perspective on the reliability of their belief-forming processes. That is, clairvoyance cases establish that mere reliability of the relevant belief-forming process(es) cannot be sufficient for the presence of justification. Rather, something else is required, and that 'something else' is to be construed in terms of an internalist constraint, such as an appropriate awareness of the evidence one has for the belief, or the general reliability of the relevant belief-forming process, and so forth. This insight can be put in terms of the Subject's Perspective Objection (SPO):

If the subject holding a belief isn't aware of what that belief has going for it, then she isn't aware of how its status is any different from a stray hunch or an arbitrary conviction. From that we may conclude that from her perspective it is an accident that

---

case shares with Norman the Clairvoyant, I intend the label 'clairvoyance cases' to include cases such as Mr. Truetemp.

her belief is true. And that implies that it isn't a justified belief. (Bergmann 2006, p. 12)

Externalist epistemologies are particularly vulnerable to the SPO given that they typically eschew any form of an awareness requirement.<sup>78</sup> Indeed, for reliabilists, so long as the belief in question is produced by a truth-conducive process, then the subject need not have *any* sort of awareness concerning what this process is, its probability of success, and so forth. With this objection in mind, then, return to the clairvoyance case. Plausibly, what underlies the intuition that reliability is insufficient for justification here is that, from Norman's perspective, it's mere accident that his belief that the President is in New York is true. Indeed, Bonjour seems to say as much in his own analysis of the case:

From [Norman's] standpoint, there is apparently no way in which he *could* know the President's whereabouts. Why then does he continue to maintain the belief that the President is in New York City? Why is not the mere fact that there is no way, as far as he knows or believes, for him to have obtained this information a sufficient reason for classifying this belief as an unfounded hunch and ceasing to accept it? And if Norman does not do this, is he not thereby being epistemically irrational and irresponsible? (1980, p. 62 – 63)

Bonjour takes this (i.e. Norman's irresponsibility were he to maintain a belief that, from his perspective, is unfounded) to establish that the belief in question cannot be even defeasibly justified.<sup>79</sup> From the internalist perspective, then, it looks like clairvoyance cases give us both a negative claim – i.e. that reliability of the relevant belief-forming process is insufficient for justification – and a positive claim – i.e. that the subject must have an appropriate awareness of the evidence she has for her belief.

Externalists, of course, reject the idea that clairvoyance cases illuminate the necessity of awareness or appropriate perspective for epistemic justification. Instead, some have suggested that we can make sense of the insufficiency of mere reliability in externalist terms.

---

<sup>78</sup> Having said that, Bergmann himself levels the SPO against internalist views. I also raise the SPO against one interpretation of Chudnoff's presentationalism, which is very plausibly internalist, in Harrison (2019).

<sup>79</sup> Although, see Bernecker (2008) for the claim that Bonjour's internalist analysis of the clairvoyance case objectionably "presupposes a bias against clairvoyance" (p. 166).

Greco (1999), for example, argues that “strange and fleeting processes” (p. 285) like the process of clairvoyance, though reliable, do not give rise to justified belief. Roughly speaking, this is because justification requires that the belief in question arises via a stable and reliable cognitive disposition integrated within the subject’s cognitive character. Greco (2003) discusses the clairvoyance case and specifies why clairvoyance is plausibly *not* a cognitively integrated process as follows:

If [Norman’s] power of clairvoyance is not well-integrated with the remainder of [his] cognition, then it will not count as cognitive character and will therefore not count as a source of knowledge according to agent reliabilism. As we have noted, one aspect of cognitive integration concerns the range of outputs – if the clairvoyant beliefs are few and far between, and if they have little relation to other beliefs in the system, then the power of clairvoyance is less well integrated on that account. Another aspect of integration... is sensitivity to counterevidence. If clairvoyant beliefs are insensitive to reasons that count against them, then this too speaks against cognitive integration. (p. 475).<sup>80</sup>

For Greco, then, acquiring valuable epistemic goods such as knowledge and justification requires that the processes through which we arrive at these goods are sufficiently entrenched and integrated parts of the subject’s cognitive character, such that clairvoyance, as it is typically described, is ruled out as being capable of performing such an epistemic role. We’ll return to the details of Greco’s proposal in §5.3.2, but, for now, note that this diagnosis of the clairvoyant’s failing is externalist. As such, Greco has a means of explaining the insufficiency of mere reliability without committing to the positive claim that something like an appropriate perspective or awareness is required for justification.

So, putting all of this together, clairvoyance cases can be interpreted as involving a negative claim that reliability of the relevant-belief-forming process is insufficient for justification, and a positive claim that something else is required for justification, and this ‘something else’ can be construed in internalist or externalist terms. Clairvoyance cases, then, bolstered and underwritten by the putative necessity of something like an appropriate perspective on one’s

---

<sup>80</sup> Note, however, that Greco is open to the possibility of clairvoyance constituting a source of knowledge, so long as that power is sufficiently well-integrated into the subject’s cognitive architecture.

belief-forming process or cognitive integration of the relevant belief-forming process, purport to establish the falsity of brute reliabilist theories of epistemic justification.

Finally, it's important to note that challenges of this sort have struck a chord with epistemologists. Indeed, many sympathetic to externalism take clairvoyance cases to be successful in establishing that reliability is insufficient for justification, and thereby cite the failure of generic reliabilism at the hands of these cases as motivating the move to an alternative theory which putatively avoids the pitfall. Comesaña (2010), for instance, argues that clairvoyance cases demonstrate that “reliability in itself is not sufficient for justification” (p. 590), and that these cases thereby motivate the move to a hybrid evidentialist reliabilist view, according to which justification requires that the belief accords with evidence which is, in turn, reliably connected to truth. Lyons (2009) also recognises the force of clairvoyance cases, writing, “[m]y own view is that the clairvoyance objection is a sound objection to SR [i.e. the generic process reliabilist view]” (p. 114). Moreover, in a similar vein to Greco, Sosa (1992) identifies clairvoyance cases – or “metaincoherence cases” (p. 81), as he calls them – as indicative of a serious fault within generic reliabilist epistemologies which at least partially motivates his own virtue reliabilist view.<sup>81</sup> Likewise, Kelp (2019) cites clairvoyance cases as one of four compelling reasons to move away from general process reliabilist views.<sup>82</sup> In light of their pervasiveness and impact throughout the literature, then, it looks like we have good reason to take clairvoyance cases seriously in an evaluation of brute reliabilism's prospects.<sup>83</sup>

Now, let us set aside the question as to whether generic reliabilism is in fact undermined by clairvoyance-style challenges, and turn our attention back to the emotions. For now, let's focus on the negative claim putatively established by clairvoyance cases, i.e. the insufficiency of reliability for justification (we'll return to the discussion of the positive claims later on in §5.3). Could an analogous challenge be levelled against ER? To determine this, consider the following case:

---

<sup>81</sup> Sosa's virtue reliabilism is a theory of epistemic justification and knowledge that we'll go on to look at in further detail in Chapter 6 (§6.2.2).

<sup>82</sup> Kelp also cites the New Evil Demon Problem, the generality problem, and the problem of world-bound reliability as three further issues for generic reliabilism.

<sup>83</sup> I take the label of 'brute reliabilism' to refer to a reliabilist view unenhanced by additional internalist or externalist supplements.

Emotional Clairvoyance: Suppose that a subject S, under certain conditions that usually obtain, is completely reliable with respect to experiencing moral disgust towards individuals that have odious moral characters. She possesses no evidence or reasons of any kind for or against the possibility of her possessing such an emotional ability. One day, S passes a stranger on the street – somebody she’s never encountered before – and experiences a clear wave of moral revulsion with respect to the stranger. As a consequence of this, S forms the evaluative belief that the stranger has an odious moral character. In fact, this evaluative belief is true and results from her emotional experience of moral disgust under circumstances in which it is completely reliable. Is S epistemically justified in believing that the stranger has an odious moral character?

This is a close analogue of the clairvoyance case. S’s experiences of moral disgust, insofar as they are psychologically immediate to S (i.e. they do not arise as mediated via any conscious judgment she has about the stranger) and insofar as S has no awareness of nor evidence for their reliability, seem to be a kind of emotional clairvoyance. S reliably forms true evaluative beliefs on the basis of these clairvoyant emotional experiences. So, the question arises: is it plausible that S is thereby immediately and defeasibly justified in her emotion-based evaluative belief about the stranger’s character?

It seems to me that the intuitive verdict will be that S is not justified. Interestingly, we might note that this case seems similar to the initial presentation of Brady’s suspicious interviewer case considered in Chapter 2 (§2.2.1). Recall that, in this initial instantiation of the over-generalisation case, the interviewer has a ‘bad feeling’ about one of the potential job candidates (to be understood as an emotional experience of suspicion or distrust towards the candidate), and, crucially, the verdict in this case seems to be that the interviewer’s emotional experience cannot even defeasibly justify their evaluative belief that the candidate is duplicitous. Indeed, the only relevant difference between the emotional clairvoyance case and the suspicious interviewer case is that the emotional reliabilist takes the reliability of S’s belief-forming process to be sufficient for justification, whereas the emotional dogmatist takes the interviewer’s emotional seeming phenomenal character to be sufficient for justification. Plausibly, then, if we have the intuition that the suspicious interviewer’s emotional experience cannot justify their evaluative belief that the candidate is duplicitous, it strikes me as plausible that we’ll also have the intuition that S’s emotional experience of

moral disgust cannot justify their evaluative belief that the stranger has an odious moral character.<sup>84</sup> Furthermore, if we *do* have this latter intuition, then it looks like ER will turn out to be false; the reliability of forming true evaluative belief on the basis of emotional experience cannot be sufficient for justification given the intuitive verdict delivered by the emotional clairvoyance example.

However, there is an immediate problem with the emotional clairvoyance case. Recall from earlier in Chapter 4 (§4.2) that I do not take the plausibility of ER to hinge on the overall plausibility of generic reliabilism. That is, I don't mean to insist that ER is a straightforward application of generic reliabilism such that the truth of the latter entails the truth of the former. Rather, for the sake of this dissertation, I intend to evaluate ER independently of the truth or falsity of generic reliabilism. The defender of ER may identify this dialectical framing as what immunises ER from the emotional clairvoyance case above. In other words, the defender of ER might insist that the emotional clairvoyance case is irrelevant to the prospects of their view given that ER does not maintain that the reliability of *clairvoyant* belief-forming processes is sufficient for justification; it maintains that emotional reliability (i.e. the reliability of forming true evaluative beliefs on the basis of emotional experiences, where this requires the reliability of the emotional system generating true evaluative contents) is sufficient for justification. In this emotional clairvoyance case, we have reliability of a belief-forming process which is entirely attributable to a power of clairvoyance, *not* the subject's emotional system. Without the power of clairvoyance, S would not be capable of reliably experiencing fitting episodes of moral disgust towards the odiousness of strangers, given that the non-evaluative features of an individual's moral character which give rise to the evaluative property of moral odiousness or abhorrence are not properties that one can track or recognise via a passing glance on the street. If this is true, then the emotional clairvoyance case is emotional in name only, and consequently, incapable

---

<sup>84</sup> One might source further support for this intuitive verdict by appeal to a similar thought experiment proposed by Carter (2020). In Carter's Metaincoherent Emoter case (p. 1242), we're invited to consider a subject, Tim, whose emotional experiences reliably track the presence of evaluative properties, such that he is maximally reliable in forming true evaluative beliefs on the basis of his emotional experiences. Crucially, however, the source of Tim's emotional reliability is the intervention of a benevolent demon. Without the demon's interference, Tim's emotional experiences would *not* reliably track evaluative properties, such that the majority of his emotion-based evaluative beliefs would turn out to be false. Carter argues that the intuitive verdict in this case is that Tim is *not* justified, despite his emotional reliability. One might construe Carter's verdict as evidence that emotional reliability is insufficient for justification. Note, however, that Carter's case specifically targets the reliability of *emotion*, rather than the reliability of the broader belief-forming process.



of constituting any real threat to ER. At most, it might provide evidence in favour of the implausibility of generic reliabilism, but not an independently assessable emotional analogue.

This strikes me as a fair objection. Perhaps a strict emotional analogue of the clairvoyance case can be dismissed by the proponent of ER on the grounds of its lying beyond the theoretical scope of the view. However, I take it that this worry can be circumvented by constructing a similar case which does not rely on sourcing the reliability of the relevant belief-forming process in the power of clairvoyance. Accordingly, consider the following example:

Caffeine Case: Suppose that, every morning on my walk into the office, I pick up a flat white from a randomly chosen local coffee shop. Some of these shops roast their coffee in a particular way such that their flat whites have very high levels of caffeine, but I have no awareness of this fact. When I drink one of these high-caffeine coffees, I reliably experience nervousness about the events ahead in my working day; i.e. the direct cause of my nervousness about the events of my day is the high concentration of caffeine that I've inadvertently consumed that morning. Now, suppose further that most every morning in which I unknowingly drink a high-caffeine coffee, I *do* have a later event in my working day which merits nervousness, e.g. a high-stakes presentation to potential employers, a meeting with a cruel and combative colleague, a gravely serious performance review, etc. As it so happens, then, my emotional experiences of nervousness concerning the events of my working day are almost perfectly reliable, but, crucially, I undergo these experiences purely in virtue of consuming strong coffee beforehand.

What makes this case relevantly analogous to the emotional clairvoyance case is that, here, we have reliability of the relevant belief-forming process facilitated by some epistemically deviant cause, i.e. the power of clairvoyance or, in this case, an unanticipated surge of caffeine at epistemically opportune moments. Without the surge of caffeine, I would *not* experience nervousness with respect to the upcoming nervousness-meriting events of my working day (suppose that I'm ordinarily an even-tempered and well-composed person). What makes this case relevantly dissimilar to the clairvoyance case, such that it's able to circumvent the above worry, is that my belief-forming process (i.e. forming evaluative beliefs concerning the high-stakes, nervousness-meriting events of my working day on the basis of

my emotional experiences of nervousness and anxiety) is entirely an emotional one. Here, I apprehend my upcoming professional events and my emotional system reliably generates true evaluative contents, such that I reliably experience nervousness with respect to the nervousness-meriting events. The belief-forming process is emotional and reliable – it just so happens that it's always prompted by an unanticipated surge of caffeine.<sup>85</sup> It seems to me that there's no compelling way in which the objector can determine this case as being beyond the theoretical scope of ER. The reliability of this belief-forming process may be ultimately attributable to the caffeine surge, but it's not the case that the *entire* belief-forming process centres around the ingestion of caffeine in the same way that the entire belief-forming process hinges on the power of clairvoyance in the above emotional clairvoyance case. In other words, my emotional system is what's reliably tracking the worrisome-making properties of my working day (albeit facilitated by the caffeine), whereas, in the emotional clairvoyance case, it's the power of clairvoyance which allows S to reliably track the odiousness-making features of the stranger's moral character.

So, with the above objection dispelled, the crucial question arises: is it plausible to determine that I'm immediately and defeasibly justified in my evaluative beliefs concerning the high-stakes, nervousness-meriting events of my working day on the basis of my emotional experiences of nervousness? Plausibly not. Again, the *only* reason that my emotional belief-forming process is reliable is because I unknowingly ingest high levels of a chemical stimulant at epistemically opportune moments. I take it that the intuition at play in the emotional clairvoyance and suspicious interviewer cases will also find a foothold in the caffeine case, and determine that, intuitively, the reliability of the relevant emotional belief-forming process is not sufficient for the presence of justification. If this is the case, then ER is an implausible theory of emotional justification.

---

<sup>85</sup> Here, the objector may raise a generality-style worry concerning the level of detail at which we individuate this belief-forming process, i.e. one might question whether the relevant level of individuation is a belief-forming process about nervousness-meriting events which occur within the professional sphere, or whether it ought to be individuated in a coarser grain. In a similar vein to my comments in Chapter 4 (§4.2.2), I invite the reader to bracket such generality-style worries for the moment, given that my proposal in Chapter 6 will provide motivation for individuating on a fine-grain basis.

### 5.3 Modifying Emotional Reliabilism

Thus far, we've seen that the caffeine case, as an analogue to traditional clairvoyance cases levelled against generic reliabilism, appears to establish the insufficiency of emotional reliability for justified evaluative belief. ER, as it stands, seems significantly less promising than it appeared to be at the end of the previous chapter.

At this point in the dialectic, one might attempt to modify the view with the aim of preserving the core emotional reliabilist thesis while also building in the theoretical resources to help exclude problematic clairvoyance-type cases. Recall from §5.2 that the traditional clairvoyance cases levelled against reliabilism are not only taken to establish the insufficiency of reliability, but rather they are also often taken to illuminate a further necessary condition for justification. We saw that internalists typically identify this condition as requiring that the subject has an appropriate awareness or perspective on the evidence one has for their belief, whereas externalists might identify this condition as requiring that the subject's belief arises via a stable and reliable cognitive disposition integrated within the subject's cognitive character. In light of these lessons putatively provided by clairvoyance cases, Comesaña has advocated on behalf of a hybrid Evidentialist Reliabilist view which purports to satisfy the internalist diagnosis of the case, while Greco has advocated on behalf of an Agent Reliabilist view which purports to satisfy the externalist diagnosis. Plausibly, then, Comesaña's and Greco's modified reliabilist proposals are plausible candidates for the basis of an analogous modified emotional reliabilist view.

In what remains of this chapter, I will present each of these interpretive options in turn, and argue that, ultimately, neither can provide a satisfactory epistemological framework for the justificatory thesis of emotion. The shortcomings of these views, combined with the above insufficiency charge, I contend, demonstrate that reliabilism does not provide us with a plausible epistemological foundation for a theory of emotional justification.

#### 5.3.1 Evidentialist Emotional Reliabilism

Corresponding to the internalist diagnosis of clairvoyance cases, we might attempt to modify ER so as to evade the problematic implications of the caffeine case by supplementing our account with a condition which captures the necessity of appropriate perspective or

awareness. In the interest of pursuing an internalist/externalist blend of emotional reliabilism, then, it seems like a natural place to look for a framework for such a view is Comesaña's Evidentialist Reliabilism.

Recall from §5.2 that Comesaña takes clairvoyance cases to successfully establish the insufficiency of reliability, and consequently advocates on behalf of a modified reliabilist view which accommodates for the putative necessity of evidence for justified belief, i.e. a version of reliabilism supplemented by aspects of an *evidentialist* view of justification. To summarise the latter view briefly, evidentialists centre around the claim that the justificatory status of one's belief depends entirely on the evidence that one has in favour of (or against) that belief.<sup>86</sup> For the most part, evidentialists tend to construe 'evidence' in internalist terms, such that two individuals that are psychologically identical in terms of their internal states will also be epistemically identical in terms of the evidence that they possess for a given proposition.<sup>87</sup>

Blending reliabilism and evidentialism, then, Comesaña advocates on behalf of the view that a belief is justified if and only if that belief is based on evidence E and the process-type of forming a belief on the basis of evidence E is reliable.<sup>88</sup> Importantly, adopting such a view seems to give us the correct result in Bonjour's clairvoyance case. That is, because Norman's clairvoyant belief is *not* based on any evidence, nor does Norman have any evidence for the reliability of his clairvoyant power, Norman is not even defeasibly justified in believing that the President is in New York. By combining the basic reliabilist requirement (i.e. the

---

<sup>86</sup> See Conee and Feldman (1985).

<sup>87</sup> Although, note that it is possible to hold an externalist-leaning evidentialist view. One might endorse Williamson's (1997) view that "a subject's evidence consists of all and only the propositions the subject knows" (p. 717), or one might adopt a weaker version of evidential externalism identified by Silins (2005), according to which one's evidence is "what one truly and justifiably believes" (p. 378). For the purposes of the discussion at hand, however, I will focus on internalist variants of the evidentialist thesis.

<sup>88</sup> Note that Comesaña's view is explicitly put forward as a theory of doxastic, rather than propositional justification, i.e. it is primarily concerned with the question of what it is for a subject's existing belief to be justified, rather than the question of what it would take for a belief to be justified regardless as to whether or not it is held by the subject. While this thesis is primarily concerned with the question of whether emotional experience can confer propositional justification to evaluative belief, I take it that this shift in focus isn't a substantive problem for our analysis, given that Comesaña suggests that the view could be straightforwardly translated in terms of propositional justification as follows: "S is propositionally justified in believing that *p* if and only if there is some evidence *e* that S has such that the type *producing a belief that p based on e* is actually a reliable process" (p. 585). For the sake of brevity and simplicity, I set this complexity aside.

reliability of the relevant-belief forming process) with the requirement of believing on the basis of evidence, then, we get a modified reliabilist view which now has the resources to be able to disarm the clairvoyance objection.

So, if Evidentialist Reliabilism generates the right result in the clairvoyance case, perhaps an emotional analogue of the view will be capable of precluding the caffeinated emoter from claiming justification for their evaluative beliefs. Taking inspiration from this idea, then, we can build such a view as follows:

Evidentialist Emotional Reliabilism (EER): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if: (i) the contents of S's emotional experience constitutes evidence for *e*, (ii) S forms *e* on the basis of that emotional experience, and (iii) S's forming *e* on the basis of that emotional experience is a reliable process of belief-formation.

Now, it should be clear that the first of these three conditions is the interesting one for our purposes. We're taking for granted in these cases that S forms the evaluative belief on the basis of the emotional experience, and Chapter 4 (§4.3) endeavoured to show that arguing on behalf of the emotions' systematic and widespread unreliability is a difficult and theoretically costly task. As such, I take it that we can plausibly assume (ii) and (iii) obtain, and focus our attention on (i).

With that in mind, then, let us turn back to the caffeine case. Is it the case that the contents of my emotional experience of nervousness constitute evidence for my evaluative belief that my upcoming professional events are worrisome? In order to answer this question, it should be clear that more detail is required concerning how we elucidate the notion of 'evidence'.

Within the internalist sphere, one can approach the question as to what constitutes a subject's evidence in one of two broad ways. First, an evidentialist could construe evidence in accessibilist terms. Accessibilism is a brand of internalism which asserts that a subject must have reflective access to the justifier (i.e. internal state). On an accessibilist spin of an evidentialist view, then, a necessary condition on a mental state serving as evidence for a given proposition is that the subject has reflective access to the contents of this mental state. Second, and alternatively, an evidentialist could construe evidence in mentalist terms. Mentalism, in contrast, asserts only that a subject's justifiers are internal mental states; the

subject need not have reflective access to these states in order for them to count as justifiers. Now, in the interest of pursuing an internalist modification of ER (given our inspiration from the internalist diagnosis of generic reliabilism's clairvoyance problem), I take it that an accessibilist spin on evidential internalism is our most promising route to clarifying the notion of evidence, given the fact that many have raised concerns about whether mentalism can accurately and satisfactorily be viewed as an *internalist* position. Notions of awareness and accessibility are typically considered to be absolutely central to internalism, such that views which purport to be internalist while rejecting the epistemic centrality of these notions are faced with a difficult question as to what permits them to retain the internalist label.<sup>89</sup> In light of this worry, then, and for the purposes of the discussion at hand, let us assume an accessibilist internalist perspective.

With our internalist element clarified, we now have a refined question: is it the case that I have reflective access to the contents of my emotional experiences of nervousness, such that they can constitute evidence for my evaluative belief that my upcoming professional events are worrisome? This strikes me as a difficult question. Speaking broadly, it seems true that subjects very often don't have reflective access to the contents of their emotional experiences. Indeed, recall from Chapter 3 (§3.2.2) that there seem to be emotional experiences which are at least *prima facie* plausible candidates for justification-conferring states which do not bring with them awareness of the relevant non-evaluative properties of the emotion's object. One example of this was an emotional experience of awe towards a piece of artwork. This seems like a case in which the experience ought to be capable of justifying one's evaluative belief that the artwork is beautiful, but the subject of the experience lacks reflective access to the non-evaluative features of the painting that the emotional experience is tracking (e.g. the particular arrangement of brushstrokes, the spectrum of hues, etc.). What makes this case relevant for our purposes, then, is that this looks like a case in which the subject lacks reflective internal access to the contents of their emotional experience of awe. If this is true, then this instance of emotional experience will not constitute evidence for one's evaluative belief, and thereby will be incapable of performing a justificatory role under EER. As I argued in §3.2.2, this seems like the wrong

---

<sup>89</sup> See, for example, Pritchard's (2006) criticism of Conee and Feldman's mentalist evidentialism: "My question... is whether it is really so obvious, as it seemingly is to Conee and Feldman, that we would refer to such a [mentalist] view as an internalist position if it lacked the accessibilist requirement. My inclination is to think that we would not" (p. 148).

result. Constraining the justificatory abilities of emotional experience on the basis of what they do and do not bring awareness of (or, in this case, what is and what is not internally accessible to the subject) generates implausible verdicts.

The problem for EER, then, is that although the view may be able to determine that the caffeinated emoter is not justified in their emotion-based evaluative belief (given that their plausible lack of reflective access to the contents of their experience of nervousness determines that their experience cannot constitute evidence for the evaluative belief), EER will also deliver the problematic result that many emotional experiences fail to constitute evidence for evaluative beliefs, and so fail to possess justificatory power, given that individuals often lack reflective access to the contents of their emotional experiences.

In response to this argument, the defender of EER may argue that individuals *do* tend to have reflective access to the contents of their emotional experiences, and this is because these contents are better spelled out in terms of the relevant evaluative property itself, rather than the non-evaluative properties of the emotional object. On this view, then, the content of my emotional experience of awe towards the painting is just the evaluative property of awesomeness putatively instantiated by the painting, rather than the set of non-evaluative properties which gives rise to that evaluative property. Plausibly, then, I *do* have reflective access to the contents of my emotional experience in this case, such that it can constitute evidence for my evaluative belief that the painting is awesome. If individuals do typically have reflective access to their emotional contents, then those contents typically can serve as evidence for evaluative belief, such that EER won't end up yielding objectionably restrictive justificatory results.

Here's the problem with this response. Conceiving of emotional contents in this way will make EER vulnerable to the over-generalisation charge. If all that's required for the contents of an emotional experience to serve as evidence for one's evaluative belief is that the individual has reflective access to the putative presence of the evaluative property that the emotion is responding to, then the caffeinated emoter in our above case will count as justified given that, presumably, she has reflective access to the evaluative property of her working events that merits nervousness. So, all in all, if we construe the contents of emotional experience in terms of non-evaluative properties, then too few emotional experiences will constitute evidence for the relevant evaluative belief, and if we construe those contents in

terms of evaluative properties, then too many emotional experiences will constitute evidence for evaluative belief.<sup>90</sup> For this reason, then, I take it that EER is implausible.

In summary, it seems that the proponent of EER faces a dilemma in how they elucidate the internalist component of the evidentialist view. Either they spin it in accessibilist terms, in which case the view runs into difficulties concerning the scope of our reflective access to emotional contents, or they spin it in mentalist terms, in which case the view faces the worry that mentalism is not accurately viewed as an internalist position. In light of these difficulties, I take it that infusing ER with an internalist evidentialist component fails to provide a plausible account as to how emotional experience can immediately and defeasibly justify evaluative belief. We'll have to look elsewhere for a promising modification of the view.

### 5.3.2 Agent Emotional Reliabilism

Given the worries facing an internalist modification of ER, then, we might instead attempt to construct a version of ER which plugs the gap exploited by clairvoyance cases with Greco's agent reliabilist framework. Recall from §5.2 that Greco identifies cognitive integration of the relevant processes as reliabilism's key to evading the problems raised by 'strange and fleeting processes', such as the clairvoyance case. Let me begin by briefly summarising Greco's view before transposing it into an epistemology of emotional experience.

On Greco's view, we must restrict the types of reliable process that are capable of giving rise to knowledge (and justified belief) to those reliable processes that result from "stable and reliable dispositions that make up [one's] cognitive character" (1999, p. 288).<sup>91</sup> Integration of these processes into one's cognitive character involves various aspects, such as the frequency of beliefs generated by the disposition, the ways in which these output beliefs relate to and interact with the subject's other beliefs, the sensitivity of the output beliefs to defeating evidence, and so forth. In a rough summarisation, we can say that a reliable process is grounded in a disposition that is integrated into a subject's cognitive character to the extent that the doxastic outputs of the process regularly interact with other beliefs and information

---

<sup>90</sup> Note that this result mirrors the dilemma that RED faced in elucidating what a truth-maker for an evaluative proposition consists in (§3.2).

<sup>91</sup> Greco's view is typically categorised as being an instance of a virtue reliabilist view. We'll look at virtue epistemology, and virtue reliabilism, in closer detail in Chapter 6 (§6.2 & §6.2.2).



(such as defeating evidence) in the subject's psychological system. With this in mind, then, the reason that Norman's belief that the President is in New York does *not* enjoy justification, despite the fact that it was formed via a reliable process, is that the process lacks cognitive integration; Norman's clairvoyant beliefs (i.e. the output beliefs) are infrequently produced, they do not regularly interact with the other beliefs in Norman's psychological system, they fail to be sensitive to counter-evidence, and so forth. Greco's Agent Reliabilism, then, seems to correctly exclude Norman from enjoying justification for his beliefs about the President's location.

So, if appeal to cognitive integration looks promising with respect to its ability to dispel the insufficiency charge raised by clairvoyance cases, then, we can inspire and formulate an emotional analogue of the view as follows:

Agent Emotional Reliabilism (AER): S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if (i) S's forming *e* on the basis of her emotional experience is a reliable process of belief-formation, and (ii) this process is sufficiently integrated within the subject's cognitive character.

At first glance, it looks like placing the notion of cognitive integration at the epistemological centre stage captures *something* about our intuition in the caffeine case. To repeat the initial description of the case, the *only* reason that my emotional belief-forming process is reliable is because I unknowingly ingest high levels of a chemical stimulant at epistemically opportune moments; without the influence of the caffeine, I wouldn't reliably experience appropriate nervousness in response to worrisome professional events. Plausibly, what's problematic about ER's determination that the caffeinated emoter is justified is that what makes their belief-forming process reliable is not down to any skill or success of theirs. Rather, the source of their emotional reliability is down to a lucky arrangement of external factors. Requiring that the belief-forming process is sufficiently integrated within the subject's cognitive character at least seems to go some way towards recognising the epistemic importance of agent-based or agent-caused reliability, rather than incidental reliability.

The problem with AER is that, while the notion of cognitive integration may capture *something* epistemically important, it's too blunt a tool to clearly exclude the problematic

emotional cases and include the good ones. That is, it's not entirely clear to me that emotional experiences that we *don't* want to be capable of justifying evaluative belief can be ruled out on the basis of they're not being sufficiently integrated within the subject's cognitive character. Reconsider the caffeine case, and recall from §5.2 that what made this case a more compelling counterexample to ER than the straightforward emotional clairvoyance case was that the belief-forming process in the former was entirely emotional, whereas the belief-forming process in the latter was primarily one of clairvoyance. Given this, it strikes me as plausible that the belief-forming process in the caffeine case *is* sufficiently integrated within the subject's cognitive character, despite the fact that it is causally sourced in the caffeine surge. The output-beliefs of this process (i.e. my evaluative beliefs concerning the worrisome nature of my upcoming professional events) plausibly interact with and are responsive to a variety of my other beliefs within my psychological system, such as my beliefs about my workplace, the relationships I have with others in this environment, the importance of having a successful career, etc. As such, it's not at all obvious to me how AER can utilise the tool of cognitive integration to exclude the caffeinated emoter from enjoying justification for her emotion-based evaluative belief.

Indeed, these limitations of cognitive integration seems particularly salient in a revised version of the biased suspicious interviewer case discussed in Chapter 3 (§3.2.1). Recall that, here, we have a scenario in which the interviewer undergoes an experience of suspicion towards a candidate which, in fact, is the result of the interviewer's implicit social bias against women. To make this case relevant, we can stipulate further that it just so happens that the interviewer's belief-forming process is reliable; most every professional woman that the interviewer experiences suspicion and distrust towards in the workplace happens to be duplicitous, such that the interviewer's forming the relevant evaluative beliefs on the basis of their emotional experiences of suspicion is a perfectly reliable belief-forming process. Is this belief-forming process also sufficiently cognitively integrated within the interviewer's cognitive character? Plausibly, yes. The bias-produced experience of suspicion very plausibly interacts with and is responsive to a number of other beliefs within the interviewer's psychological system, such as their implicit misogynistic beliefs about the trustworthiness of career-driven woman, the superiority of male character in the workplace, and so forth. Despite the fact that the biased interviewer case meets both conditions for AER, however, it

seems implausible to allow justificatory power to their emotional experiences.<sup>92</sup> Thus, in lieu of an explanation as to how the process of forming an evaluative belief on the basis of emotional experience is sufficiently cognitively integrated in the good cases, but *not* in the bad cases, then, I take it that AER can't provide a plausible explanation as to how certain emotional experiences can immediately and defeasibly justify evaluative belief.

## 5.4 Summing Up

The purpose of this chapter has been to present what I take to be the most significant problem facing emotional reliabilism (ER), namely, the objection that the reliability of forming evaluative beliefs on the basis of one's emotional experience is insufficient for the epistemic justification of evaluative belief.

I began in §5.2 by introducing clairvoyance cases, and explaining how these cases are typically utilised in order to establish: (i) that reliability of the relevant belief-forming process is insufficient for epistemic justification, and (ii) that something else is required for justification, where this 'something else' can be construed in either internalist or externalist terms. I then constructed an analogous emotional case which aimed to demonstrate that the reliability of forming evaluative beliefs on the basis of one's emotional experience cannot be sufficient for the justification of evaluative belief.

Having levelled the objection against ER, §5.3 then considered the possibility of modifying ER in light of clairvoyance cases. In this section, I presented two alternative emotional reliabilist views corresponding to the internalist and externalist diagnoses of the clairvoyance case respectively. §5.3.1 presented the Comesaña-inspired Evidentialist Emotional Reliabilism (EER), and §5.3.2 presented the Greco-inspired Agent Emotional Reliabilism (AER). I argued that both EER and AER were unable to provide a plausible explanation as to how emotional experience can immediately and defeasibly justify evaluative belief. Given

---

<sup>92</sup> At this point, the defender of AER might object, in a similar vein to a response raised in Chapter 2 (§2.2.1) that the intuitive verdict delivered by the biased interviewer case is being hijacked by a moral intuition that the interviewer is making a *moral* mistake in harbouring and believing on the basis of such a bias. However, because the caffeine case plausibly lacks this moral element and generates the same intuitive result, I take it that we have good reason to suppose that there is *also* a relevant epistemic violation in the biased interviewer case, such that it remains relevant and problematic for the defender of AER.

the failures of ER, EER, and AER, then, I conclude that the prospects for a reliabilist approach to emotional justification appear to be bleak.

Where does this leave us? I take it that the discussions contained within Chapters 4 and 5 illuminate several interesting findings with respect to the epistemic relevance of reliability when it comes to how emotional experience can justify evaluative belief. First, although the sufficiency of emotional reliability has been undermined in this chapter, I take it that we still have good reason to think that reliability, in some sense, may be necessary for epistemic justification. Building a reliability condition into our account, i.e. emphasising a robust connection between emotional experience and evaluative truth, allowed us to both sidestep the specific kinds of over-generalisation cases levelled against dogmatist views and capture a wide variety of support from authors in the literature. Moreover, recall that the New Evil Demon Problem failed to get a foothold on ER; an emotional version of the NEDP was insufficiently powerful to establish that emotional reliability cannot be necessary for emotional justification. Finally, although the Greco-inspired AER ultimately failed to provide us with a compelling theory of emotional justification given our worries with the notion of cognitive integration, it seems plausible that turning our attention to features of the agent *herself* provides us with an interesting and promising perspective on how we might be able to build a more nuanced understanding as to how emotional experience can justify evaluative belief. Plausibly, it's honing in on the reliability of the belief-forming processes in *isolation* from the rest of the subject's epistemically relevant features that renders a view vulnerable to clairvoyance-style challenges. Turning our attention instead to the way in which this reliable belief-forming process is situated and contextualised within the subject's overall cognitive character may help us to evade these challenges, and provide a more compelling account of emotional justification.

Looking forwards, then, Chapter 6 will provide an in-depth analysis into the prospects of agent-based views of emotional justification. I believe that it is a specific view within this category that provides the most plausible framework for the justificatory thesis of emotion.

## Chapter 6

### Virtue, Character, and Competence: Agent-Based Approaches

#### 6.1. Taking Stock

Thus far in the thesis, I've identified and evaluated two general approaches to a theory of emotional justification: emotional dogmatism and emotional reliabilism. Broadly speaking, we've seen that the downfall of dogmatism consists in the view's inability to provide a plausible phenomenology-based explanation as to how we can suitably restrict the view so as to exclude over-generalisation cases from enjoying epistemic justification, while reliabilism is rendered untenable by a vulnerability to an emotional analogue of the clairvoyance case, i.e. a case which establishes the insufficiency of emotional reliability for epistemic justification.

Before turning our attention to the third and final category of views, it will be instructive to first reflect upon the various desiderata for a plausible account of emotional justification raised by the vices of dogmatism and reliabilism. I take it that there are three desirable criteria that an account of emotional justification ought to be able to meet in order to constitute an improvement over the previous views. The first desideratum is that the theory must be suitably restrictive, such that it offers a principled reason as to why the biased interviewer's baseless emotional experience of suspicion towards the candidate cannot immediately and defeasibly justify their evaluative belief that the candidate is untrustworthy, for example. Second, on the other side of the coin, the view must not veer *too* far on the side of restraint, such that emotional experiences that we would ordinarily consider to be valuable justificatory resources do not end up being excluded on spurious grounds. Similarly, recall that for the justificatory thesis of emotion to be of particular epistemological importance, it must promise a sufficiently widespread degree of emotionally justified evaluative belief, such that emotional experience can also serve as the experiential bedrock for further epistemic goods, such as evaluative knowledge and understanding. Third, bearing in mind the clairvoyance-type challenges raised against emotional reliabilism, the view must recognise and provide an explanation of the necessity but insufficiency of emotional reliability. That is, while reliabilism's focus on a robust connection between experiential justification and truth was a welcome development in some respects, emotional reliability alone cannot be enough to guarantee positive justificatory status for the relevant evaluative beliefs.

So, where do we look for a view which can accommodate these desiderata? The final category of views that will be examined in this dissertation are theories of emotional justification grounded in a virtue epistemological framework; what we might call *agent-based* views of emotional justification. These are views which prompt a shift in focus from assessing features of the emotional experience and the emotion-based evaluative beliefs to assessing features of the emoter *herself*, such as the intellectual character traits or epistemic competences that led to the formation of her emotional experience and the corresponding evaluative belief. On agent-based views, it doesn't make sense to epistemically evaluate the emotional experience and the corresponding belief independently of the way in which they were formed by the emoter. Instead, we must direct our attention to the etiology of these experiences, and base our epistemic assessment on the features of the emoter that give rise to the emotional experience in question. Here, I argue that an agent-based view is our best bet for a plausible account of emotional justification. More specifically, I will argue that an account grounded in the possession of learned emotional competences both satisfies the above desiderata, and yields further philosophical advantages which make it, in my view, the most promising view of emotional justification.

The structure of this chapter is as follows. In §6.2, I provide a brief introduction to the topic of virtue epistemology, and set out the two broad virtue epistemological approaches: virtue responsibilism (§6.2.1) and virtue reliabilism (§6.2.2). Here, I suggest that virtue responsibilism is unlikely to be able to provide a plausible framework for a theory of emotional justification, and that we ought to focus our efforts on building a virtue reliabilist theory instead. On that note, §6.3 then presents what I take to be the most plausible instantiation of such a view available on the philosophical market, namely, Carter's competence-based view of emotional justification (EC). This section clarifies the account and highlights its advantages.

In spite of these advantages, however, §6.4 then argues that EC is in need of further development. More specifically, I identify an area of ambiguity within the account which requires clarification. Namely, given that the possession of generative emotional competence requires that the subject is likely to succeed in experiencing an emotion that is appropriately responsive to the presence of an evaluative property, I argue that the scope of this success requirement is unclear. As it stands, EC is ambiguous between the following three options: (i)

the likelihood of success is calculated by taking into account *all* of the subject's emotional experiences, (ii) the likelihood of success is calculated by taking into account *only* instances of the given emotion-type, (iii) the likelihood of success is calculated by taking into account *only* instances of the given emotion-type *elicited by* a particular scenario. The upshot of this ambiguity is twofold. First, without a clear explication of the success requirement, the view fails to provide a fully informative account as to what is required for the possession of emotional competence, and, consequently, what is required for emotional justification. Second, and relatedly, if it is the case that calculations of likely success must only take into account instances of the given emotion-type (or emotion-type + particular eliciting scenarios), then it would seem to follow that we can have emotional competences localised to particular emotion-types (or emotion-types + particular eliciting scenarios). If this is true, then it seems to follow that it is possible for *some* of a subject's emotional experiences to bear justificatory power with respect to their evaluative beliefs, but not others. It is for these reasons that I take clarifying this ambiguity to be essential. I then dismiss option (i) as being insufficiently nuanced to accommodate for the complexity of our emotional skills and profiles, and propose that we turn our attention to the localised variants as a more plausible starting point.

Accordingly, §6.5 then argues that EC's success requirement on emotional competences ought to be construed in terms of option (iii), i.e. as localised to particular emotion-types experienced in response to particular eliciting scenarios. This is because the most plausible explanation of the development of generative emotional competence, I contend, is that individuals undergo processes of emotional learning with respect to the presence of particular evaluative properties in particular eliciting scenarios. My view is that it is when and only when an emotional experience arises from a competence forged by these learning processes that it is capable of immediately and defeasibly justifying evaluative belief. This refined version of EC is what I refer to as the Learned Emotional Competence (LEC) theory of emotional justification. I then strive to bolster the prospects of LEC by expanding upon the processes and significance of emotional learning (§6.5.1), and highlighting a unique selling point of LEC, namely, its ability to illuminate and explain a popular insight from feminist philosophy (§6.5.2). I then close in §6.6 by summarising the virtues of LEC and explaining how it meets the desiderata identified for a compelling theory of emotional justification. Finally, I provide concluding remarks in §6.7.

## 6.2. Virtue Epistemology and Emotional Justification

As noted above, virtue epistemologists are concerned with shifting the locus of epistemic assessment from the qualities of the belief in question to the qualities of the believer. That is, it is features of the epistemic agent *herself* that matter for the belief's epistemic status, such as her epistemic virtues, skills, or competences, rather than questions of whether the belief was produced via a reliable process, or whether it was based on an experience with the appropriate phenomenal character, and so forth. Looking at justified belief through a virtue epistemological lens, then, we arrive at the rough idea that an agent's belief that *p* is justified if and only if that belief is the result of the agent's exercise of the relevant intellectual capacity or feature.

With respect to answering the question of what *kind* of agential feature is relevant for epistemic assessment, virtue epistemologists have been traditionally divided into one of two broad camps. Roughly speaking, virtue *responsibilists* attribute epistemic significance to an agent's intellectual and epistemic character traits, such as her curiosity, open-mindedness, intellectual courage, and so forth, whereas virtue *reliabilists* attribute significance to an agent's reliable cognitive faculties and capacities, such as her perception, introspection, intuition, and so forth. On the question of what makes a belief a *justified* belief, then, virtue responsibilists inclined to answer such a question will argue that a justified belief is a belief that arises via the exercise of the agent's intellectual character traits, while virtue reliabilists will insist that a justified belief is a belief that arises via the exercise of her reliable cognitive faculties. Thus, while both of these views identify the seat of a belief's positive justificatory status as embedded in the epistemically relevant features of the believer, they differ on what these relevant features are.

In what remains of this section, I will sketch out both virtue responsibilism and virtue reliabilism in finer detail, and set out how we might figure emotional experience into each of these justificatory pictures. More specifically, I will argue that virtue responsibilism is unlikely to be able to provide a satisfactory model for a justificatory thesis of emotion, given its focus on intellectual character traits and virtuous motivation or inquiry. Instead, I suggest that we ought to look to virtue reliabilism for a plausible agent-based view of emotional justification.



### 6.2.1. Virtue Responsibilism

Let us first consider the prominent virtue responsibilist view elucidated by Zagzebski. On the topic of justified belief, Zagzebski (1996) suggests the following:

A justified belief is what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, might believe in like circumstances. (p. 214)

First, and centrally, we have the notion of intellectual virtue. For Zagzebski, intellectual virtues have both a *motivational* and a *reliability* component. Take, as an illustrative example, the intellectual virtue of open-mindedness. On Zagzebski's view, open-mindedness involves an affective motive towards a particular valuable end, i.e. a love or care for considering appropriate alternative ideas. Crucially, however, it also involves a reliability component insofar as the open-minded agent will be reliably successful in considering appropriate alternative ideas. To generalise, the motivational component of each intellectual virtue is derivative from a general care for, or love of, truth, and it is relativised to the particular trait in question, e.g. curiosity involves the motive of care to discover new ideas, intellectual courage involves the motive of care to persist in holding beliefs and conducting epistemic inquiries even when under pressure to discard them, and so forth. Moreover, each motivational component is matched by a reliable success condition, such that the intellectually virtuous agent will have the disposition to be appropriately motivated towards suitable ends and will be reliably successful at achieving those ends.

The second feature of Zagzebski's account of epistemic justification requires that the subject has the relevant understanding of their cognitive situation. Put simply, what this requirement amounts to is the subject possessing the relevant descriptive facts pertaining to their epistemic situation and the formation of the relevant belief, such as relevant facts about the contents of the belief, their evidence and their surroundings, and so forth. In summary, then, Zagzebski provides a picture of justified belief in which a belief is justified if that belief *would be* formed by an agent who (i) is motivated by intellectual virtue, and (ii) has the relevant understanding of her epistemic situation pertinent to the formation of her belief.

So, how might emotional experience figure into this picture? The most obvious opening seems to be the affective element of the motivational component of intellectual virtues. Indeed, Zagzebski explicitly leaves epistemic space for the emotions in her description of the relevance and nature of motivation:

What is important is just the idea that beliefs can be and perhaps typically are motivated, and that the motive can affect the evaluation of the belief in a way that is analogous to the way the motive can affect the evaluation of an overt act. What I mean by a motive is an affective state that initiates and directs action. In my theory of emotion, a motive is an emotion that is operating to produce action. (2003, p. 17)

In light of these comments, then, perhaps, we can construct a virtue responsibilist account of emotional justification along the following lines:

S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if that experience at least partially constitutes an intellectual virtue.

There are, however, some issues with endorsing such a view. First, it's not at all clear that our various everyday emotional experiences will be capable of justifying evaluative belief along Zagzebski's lines. It is not obvious how, say, fear of an approaching snake or disgust towards spoiled milk can figure into an intellectual virtue in the relevant way so as to be bestowed with justificatory power. These emotional experiences do not obviously seem to be tied up with any valuable intellectual motivation (to uncover truth, to consider new ideas, etc.) such that they can be considered as partly constitutive of intellectual virtue.<sup>93</sup> If it's *only* emotional experiences that neatly figure into epistemic inquiries of this sort that are capable of bearing justificatory power, then one might worry that this version of the justificatory thesis fails to satisfy the above desideratum of delivering a substantive enough epistemic yield.

---

<sup>93</sup> Perhaps it might be argued that these emotional experiences *are* tied up with an intellectually valuable motivation, i.e. to discover truth (where we construe this motivation in a *very* broad sense so as to include fairly trivial or banal truths, such as the fearsomeness of nearby threats or the disgustingness of spoiled foods). If one adopts this line of argument, however, it seems to me that most every type of emotional experience could be described in such a way insofar as emotional experiences putatively purport to represent world accurately, such that most every emotional experience would count as partially constituting an intellectual virtue. If this were the case, of course, the view would over-generalise.

Second, and relatedly, a view of this sort only seems to offer emotional experience a backseat role in securing epistemic goods in a way which clashes with the original motivations of the justificatory thesis. That is, it's not *really* the emotional experience doing the fundamental epistemic work here; the emotional experience is only capable of performing an epistemic role insofar as it figures into the broader intellectual virtue. It's the *latter* that's particularly epistemically significant, and this seems at odds with the motivations that make the justificatory thesis of emotion philosophically interesting, i.e. that emotional experience acts as a simple and naturalised means of obtaining justification for evaluative belief, and, in turn, aids in the acquisition of further epistemic goods. Indeed, on this note, many of Zagzebski's critics object to the idea that the motivational component of intellectual virtue is fundamentally affective or emotional.<sup>94</sup> In response to criticisms along these lines, Zagzebski concedes that "the formal structure of the theory does not require that motives be emotions" (2000, p. 210). Furthermore, it's not clear how such a view could explain the *immediacy* of the justification conferred by emotional experience. If emotional experience can only justify evaluative belief when it constitutes an intellectual virtue, and possessing intellectual virtue requires an understanding of one's cognitive situation, then one might worry that the justificatory ability of emotional experience hinges on the possession of that non-emotional cognitive understanding in such a way that undermines the former's capacity to justify evaluative belief independently of any connections to the subject's existing justified beliefs. Given these concerns, then, it looks like a justificatory thesis of emotion built on the foundations of Zagzebski's virtue responsibilism has notable shortcomings.

Importantly, though, it's not just Zagzebski's instantiation of the responsibilist view that gives rise to worries like these for a responsibilist account of emotional justification. Consider, on the other hand, Hookway's comments on the importance of epistemic and intellectual character when it comes to determining the justificatory status of one's belief:

[A]t best we might say that someone is justified in believing something if their belief issues from responsible and virtuous inquiry. Justified beliefs are those that issue from the responsible inquiries of virtuous inquirers. (1994, p. 225)

---

<sup>94</sup> See Alston (2000), p. 185.

Hookway's account differs from Zagzebski's insofar as the central focus is on the importance of virtuous inquiry, rather than acting and believing in accordance with intellectual virtue (conceived as consisting in the motivational and reliability components). For Hookway, responsible and virtuous inquiry is a deliberative reasoning process unconsciously guided by cognitive character traits (i.e. intellectual virtues) which are embedded in one's character through experience and practice. On this view, the intellectually virtuous agent will *naturally* engage in virtuous inquiry which manifests their entrenched traits of open-mindedness, curiosity, intellectual courage, and so forth. It is the beliefs arrived at through this process that count as being epistemically justified.

How, then, might emotional experience enter the picture here? Hookway is receptive to the suggestion that emotions provide a positive contribution to effective inquiry. On the topic of the role of emotional experience in scientific deliberation, Hookway (2002) suggests that the emotions play a number of significant mechanical roles in regulating this inquiry, such as driving and halting inquiry through the initiation and cessation of epistemic anxiety respectively, acting as "cognitive shortcuts" (p. 257) which eliminate the need for burdensome and complex assessments of evidence, and so forth. Again, however, it's not obvious that this claim – that emotional experience is necessary for virtuous inquiry, which, in turn, is necessary for justified belief – can get us to the view that emotional experience *itself* immediately and defeasibly justifies evaluative belief. Perhaps the most that can be said about the emotions here is that they *figure into* the system which produces justified belief. Thus, it looks like Hookway's virtue responsibilism, as a potential basis for a justificatory thesis of emotion, suffers similar shortcomings to Zagzebski's view.

Taking all of this into account, it seems to be the case that virtue responsibilist views lack the resources to be able to easily offer up an account of immediate experiential justification. This isn't particularly surprising, given the general virtue responsibilist motivation to shift the focus of traditional epistemic theorising from notions of 'justification' and 'knowledge' to more holistic notions of 'intellectual character' and 'virtue'. Take Hookway's claim that, if his responsibilist account is correct, "a general notion of 'justification' may have relatively little work to do in a developed account of epistemic evaluation" (1994, p. 225), or Zagzebski's intent to offer a virtue responsibilist view that gives the concept of epistemic justification "a more modest role in epistemic evaluation" (2000, p. 213). Epistemic justification just *isn't* of fundamental interest to virtue responsibilists, so it is to be expected

that these views will not lend themselves particularly well to forming a basis for a corresponding theory of emotional justification. In other words, virtue responsibilism and the justificatory thesis of emotion seem to be pursuing separate epistemological projects, such that a smooth blend of the two seems unlikely.

In sum, if the defender of the justificatory thesis of emotion wishes to employ a virtue responsibilist framework to spell out their view, then it looks like they'll have to engage in some tricky theorising to reconcile the aims of virtue responsibilism with the justificatory thesis of emotion, and explain both: (i) how our emotional experiences figure into the possession of intellectual virtue, and (ii) how our everyday emotional experiences, which are *not* obviously tied to epistemic inquiry, can be capable of bearing immediate justificatory power. For these reasons, then, perhaps we ought to look elsewhere in the virtue epistemological sphere for a plausible account as to how and when emotional experience can immediately and defeasibly justify evaluative belief.

### **6.2.2. Virtue Reliabilism**

Let us now turn to virtue reliabilism as an alternative epistemological framework for the justificatory thesis of emotion. Recall that virtue reliabilists determine that intellectual virtue, i.e. what makes beliefs epistemically valuable, amounts to the subject's possession of certain faculties which reliably give rise to true beliefs. Paradigmatic intellectual virtues, on such a view, are typically identified as faculties such as perception, intuition, memory, and the like. Where virtue reliabilists tend to differ is at the level of detailing the exact nature of these intellectual virtues, and their relationship to the believer.

Now, interestingly, we've already come across an example of a virtue reliabilist view in our discussion. Recall Greco's agent reliabilist view discussed in Chapter 5 (§5.3.2), according to which the reliable belief-forming processes that are capable of giving rise to knowledge (and justified belief) are those reliable processes that result from "stable and reliable dispositions that make up [one's] cognitive character" (1999, p. 288). For Greco, these integrated dispositions constitute cognitive (or intellectual) virtues. Recall also, however, that Greco's view encountered issues when transposed to emotional experience given worries concerning the notion of cognitive integration. For the purposes of this chapter, then, I will focus solely on an alternative virtue reliabilist view, namely, Sosa's instantiation of the account, given

both the impact and influence that Sosa's view has had within the literature and the results of our prior discussion of Greco's view. In what follows, I will provide a brief overview of Sosa's broad epistemological picture, before highlighting the place of epistemic justification within it, and explaining why such an account is a more promising epistemological foundation for the justificatory thesis of emotion.

The starting point for elucidating Sosa's view concerns his claim that epistemic normativity is a kind of performance normativity.<sup>95</sup> Roughly speaking, on Sosa's view, beliefs are types of performances aimed at truth, and, as performances, are assessable on the following three levels. First, performances can be *accurate*, i.e. whether the performance is successful in its aim. Second, performances can be *adroit*, i.e. whether the performance manifests a particular competence of the performer. Third, performances can be *apt*, i.e. whether they are accurate *because* they are adroit. To summarise a complex view briefly, Sosa argues that an *apt* belief constitutes an instance of knowledge.<sup>96</sup> We secure valuable epistemic goods, therefore, when our epistemic performances manifest our competences in the relevant way, where these competences are to be understood as dispositions to reliably succeed in our particular performances.

So, if knowledge amounts to apt belief, where does epistemic justification enter the picture? In Sosa's early work, he provides the following comments on epistemic justification:

Here primary justification would apply to intellectual virtues, to stable dispositions for belief acquisition, through their greater contribution toward getting us to the truth. Secondary justification would then attach to particular beliefs in virtue of their source in intellectual virtues or other such justified dispositions. (1980, p. 23)

---

<sup>95</sup> Indeed, one can find echoes of Sosa's performance-based virtue epistemology in the work of other notable virtue epistemologists. See, for example, Pritchard (2009).

<sup>96</sup> Technically speaking, Sosa identifies apt belief as an instance of *animal* knowledge, which is presented in contrast to the more demanding *reflective* knowledge. The latter requires that the belief is apt *and* that the believer has the appropriate first-person perspective on the aptness of her belief. Given our focus on epistemic justification, rather than knowledge, I won't expand on this detail here, but note that this additional perspective requirement on the possession of reflective knowledge is useful insofar as it makes Sosa's account capable of determining that Norman's clairvoyance-produced belief about the President's location does not count as an instance of reflective knowledge.

What Sosa refers to here as ‘secondary justification’ is the kind of epistemic justification that we’re interested in here; the kind that attaches to particular beliefs when they arise from intellectual virtues, i.e. dispositions to form true beliefs. In a similar vein, we can observe a similar point put differently in Sosa’s later comments:

A belief may similarly hit the mark of truth unaided by luck, and may also fit within the believer's wider body of beliefs. And we can then evaluate it as “epistemically justified,” in one or another sense: “competently adroit” perhaps (or reliably based, or counterfactually safe, etc.)... (2009a, p. 114)

Here, Sosa provides an analysis of epistemically justified belief in terms of that belief being “competently adroit”, i.e. manifesting subject’s competence to form true beliefs. Taking these comments into account, we see that, on Sosa’s performance-based virtue reliabilism, a belief’s being epistemically justified amounts to that belief being produced via the subject’s epistemic competence, i.e. their disposition to reliably succeed in forming true beliefs. Of course, reliable success does not entail infallible success. All that is required for a belief to be justified, on Sosa’s view, is that the subject’s formation of the belief manifests the competence to reliably *enough* form true belief. The possession and manifestation of particular competences are thus the key to this virtue reliabilist picture of epistemic justification.

What are competences, then? On Sosa’s view, competences are “dispositions of an agent to perform well” (2010, p. 465) , and can be analysed in terms of their possessing a ‘triple S’ structure. To illustrate, take an individual with perceptual competence. The possession of a competence involves three constitutive components. First, there’s the *skill* (or *seat*), which constitutes the innermost competence as the disposition to perform the activity successfully, e.g. the skill of perceiving the sensible world veridically. Second, the disposition requires that the subject be in the proper *shape*. That is, she’s in the relevant internal condition to be able to execute her skill, i.e. she’s not under the influence of a hallucinogen. Third, and finally, the disposition requires that the subject be properly *situated*. That is, her external and environmental conditions must be such that she is capable of manifesting the skill, e.g. she’s not in poor lighting conditions.

Now, briefly, there are a couple of things to note here. First, note that the subject retains possession of the skill even when she is not properly shaped and situated. That is, if I'm reliably successful when driving, then I still possess the skill to drive even if I'm submerged several hundred metres underwater and have recently ingested a hallucinogenic, but, crucially, I don't possess the competence if I can't *actually* drive successfully because I'm improperly shaped and situated. Second, generally speaking, identification of the relevant shape/situation pairs which correspond to the skill (which, constitutively, give rise to the relevant competence) is dependent on the background norms which govern and contextualise the performance of the act in question. The relevant shape/situation pair for assessing whether an individual has an archery competence, for instance, will be indexed to the conditions in which good archery performances are valued (i.e. in non-torrential weather conditions, when the archers are not under the influence of alcohol, etc.).

Taking all of this into account, we can thereby test for an agent's possession of complete perceptual competence with the following trigger-manifestation conditional: if our subject is *likely* to succeed in undergoing a veridical perceptual experience, given that she's properly situated and in proper shape, then she has perceptual competence. Thus, forming a perceptual belief as a result of this competence would constitute an adroit performance, and, thereby, be evaluated as an epistemically justified belief. This is the essence of Sosa's performance-based virtue reliabilist view.

With all of this in mind, then, we can determine straight away why Sosa's virtue reliabilism is better placed than virtue responsibilism for our investigation into an agent-based view of emotional justification. First, given that virtue reliabilism's focus is on particular competences and faculties, such as perception and intuition, there's a sense in which we can straightforwardly 'plug-in' a subject's emotional system as constituting one of these faculties. This makes virtue reliabilism much more amenable to lending itself to accounts of immediate experiential justification than virtue responsibilism. Second, and relatedly, we can allow the emotions a central epistemic role without necessitating that they're intricately bound up with intellectually virtuous character traits, or virtuous inquiry, and so forth. Moreover, as we'll see in what follows, marrying the justificatory thesis of emotion with a virtue reliabilist framework meets the various desiderata specified for a plausible and compelling account of emotional justification. Without further ado, then, the next section will present a view which attempts to explain how emotional experience can immediately and defeasibly justify



evaluative belief with reference to this virtue reliabilist framework. This, I take it, is the best starting point for a plausible agent-based theory of emotional justification.

### 6.3. The Emotional Competence View

Fortuitously for our purposes, a view of emotional justification grounded in Sosa's virtue reliabilist framework has recently been developed by Carter (2020). In the interest of unifying a competence-based virtue epistemology with a justificatory thesis of emotion, Carter proposes the following:

Emotional Competence View (EC): S's emotional experience immediately and defeasibly justifies her evaluative belief *e* if and only if (i) the emotional experience is formed aptly (i.e. manifests S's generative emotional competence) and (ii) S forms *e* via the exercise of her doxastic emotional competence.<sup>97</sup>

To illustrate, reconsider an experience of anger in response to an offhand remark made by a colleague. According to EC, this experience is capable of justifying my evaluative belief that the remark is offensive if and only if the following two conditions are met. First, the experience of anger must be formed aptly, i.e., the experience must manifest a *generative emotional competence* to experience emotions which are appropriately responsive to the presence of an evaluative property. That is, roughly speaking, one has generative emotional competence insofar as one possesses the competence of reliably experiencing anger towards the offensive, shame towards the shameful, and so forth. Crucially, however, mere generative emotional competence is not sufficient for the positive justificatory status of the relevant belief. That is, it's not enough to have a competence to *experience* the appropriate emotions, but, rather, one must also be competent at *forming* the appropriate evaluative belief on the basis of their apt emotional experiences; at transitioning from an apt emotional experience to the relevant evaluative belief. This leads us to the second condition: my evaluative belief that the remark was offensive must be generated via a *doxastic emotional competence* to form the appropriate evaluative belief on the basis of the relevant emotional experience. The apt emotional experience can only justify the evaluative belief when there's an appropriate causal

---

<sup>97</sup> See Carter (2020), p. 21. The 'Emotional Competence View (EC)' is my label for the view.

connection between the two. Thus, for Carter, it is *only* when these emotional skills are working in tandem that an agent's emotional experiences can justify her evaluative beliefs.<sup>98</sup>

How, exactly, should we understand generative and doxastic emotional competence? Recall Sosa's triple-S structure of competences as elucidated above. Adopting that model, we can test for the possession of emotional competences in terms of the following trigger-manifestation conditionals:

Generative Emotional Competence: if a subject S is likely to succeed in experiencing an emotion appropriately responsive to the presence of an evaluative property, given that she's properly shaped and situated, then she has generative emotional competence.

Doxastic Emotional Competence: if a subject S is likely to succeed in forming the appropriate evaluative belief on the basis of her apt emotional experiences, given that she's properly shaped and situated, then she has doxastic emotional competence.

Now, a question may arise concerning how we identify the relevant shape/situation pairs for the domain of generative and doxastic emotional competence. That is, there is a question as to what constitutes proper shape and situation for experiencing appropriate emotional experiences and forming the relevant evaluative beliefs on the basis of these experiences. Recall that the relevant shape/situation pairs vary across different domains of performance, and are indexed to the norms in which good instances of that performance are valued. For the purposes of this discussion, I take it that it will suffice to make common sense assumptions about proper shape and situation required for generative and doxastic emotional competence. That is, she is not suffering from mental illness or psychological disorder, she is not under inordinate pressure in a life-or-death situation, and so forth.

---

<sup>98</sup> One interesting thing to note about the relationship between generative and doxastic emotional competence, as Carter spells it out, is that it doesn't seem to be possible to meet condition (ii) of EC without first meeting condition (i). This is because exercising one's doxastic emotional competence in a way which meets condition (ii) involves, as Carter puts it, "taking not just any emotion, but an apt emotion at face value" (p. 1248). In other words, if S forms *e* via the exercise of her doxastic emotional competence, she must be transitioning from an emotional experience which manifests her generative emotional competence to the relevant evaluative belief. Thus, it looks like (i) must be met before (ii) can be met.

Now, having set out the view, we're now in a position to be able to evaluate EC with respect to the initial desiderata set out in §6.1. Recall that an essential criterion identified for a plausible justificatory thesis of emotion is that it is suitably restrictive such that emotional experiences like those of the biased interviewer don't turn out to be capable of immediately and defeasibly justifying evaluative belief. Recall that, in this example, the interviewer has an emotional experience of distrust or suspicion towards the candidate which in fact arises from their illicit bias against members of the candidate's social group. Now, it should be clear that EC will not bestow this emotional experience with justificatory power. Neither the interviewer's emotional experience *nor* their corresponding evaluative belief that the candidate is duplicitous manifest the interviewer's generative or doxastic emotional competences respectively. A social bias will very likely not constitute an emotional competence given that biases against particular social groups cannot be plausibly construed as dispositions to perform well within the domain of experiencing emotions appropriately responsive to the relevant property. Indeed, it seems as though biases against particular social groups can be identified as dispositions to experience emotions which are inappropriately responsive to the presence of irrelevant non-evaluative properties, e.g. an individual's appearance properties. On the other hand, recall that a plausible theory of emotional justification must not veer too far in the direction of restraint, and rule out certain emotional experiences from bearing justificatory power on spurious grounds. On the face of it, EC seems to meet this requirement. For EC, *any* emotional experience *could* be capable of bearing justificatory power, so long as the competence requirements are met. Importantly, it doesn't seem *unreasonably* exclusionary to determine that a subject's emotional experience is incapable of bearing justificatory power when that subject lacks the relevant emotional competences. So far, then, EC seems to be in significantly better shape than the dogmatist approaches to emotional justification considered in Chapters 2 and 3.

What's interesting to consider now is whether EC fares better than the reliabilist approach. Recall that our third desiderata for a plausible view was an ability to explain and account for the insufficiency of mere emotional reliability for justification. This insufficiency was established in Chapter 5 (§5.2) by appeal to a counterexample relevantly analogous to traditional clairvoyance cases levelled against generic reliabilism. In this counterexample, we considered a subject who is typically reliable with respect to forming evaluative beliefs concerning the anxiety-meriting events of their working day on the basis of their anxiety experiences, and these emotional experiences typically generate reliable evaluative contents,

but the reason for this is that every time the subject does have an anxiety-meriting working day, they happen to have a strong cup of coffee that morning, and the surge of caffeine is what explains their anxiety. I argued that this looks like a case in which emotional reliabilism will generate a false positive. EC, on the other hand, looks particularly well-equipped for explaining why the subject's anxiety-based evaluative beliefs are not justified. Simply put, the subject's experiences of anxiety cannot be capable of bearing justificatory power because they do not manifest a generative emotional competence. By stipulation, removing the influence of caffeine from the case also removes their likelihood of succeeding in experiencing anxiety appropriately responsive to the anxiety-making features of their working day. If this is the case, then EC can determine that the subject's emotional experiences of anxiety are not capable of immediately and defeasibly justifying the relevant evaluative beliefs in virtue of the fact that they do not manifest or arise via the required emotional competence. Thus, it looks like EC is well-equipped for recognising the necessity but insufficiency of mere emotional reliability for justification; reliability and likely success are epistemically essential for EC, but, crucially, this reliability *must* result from the subject's generative and doxastic competences.

So, thus far, the prospects of EC appear to be promising. Insofar as the view is capable of evading the challenges which troubled dogmatist and reliabilist approaches, it looks like EC is particularly well-placed for a plausible view as to how emotional experience can immediately and defeasibly justify evaluative belief. However, in what follows, I will argue that EC, in fact, requires further development and clarification in light of an ambiguity in the scope of emotional competences. In light of this ambiguity, I will then propose a crystallised version of the Emotional Competence view which, I believe, is a stronger contender for a compelling view of emotional justification.

#### **6.4. Objection: An Ambiguity in the Scope of Emotional Competences**

Put simply, the worry is that EC's overall plausibility will hinge on how, exactly, we spell out the scope of a subject's emotional competences. Reconsider the trigger-manifestation conditional for generative emotional competence:

If a subject S is likely to succeed in experiencing an emotion appropriately responsive to the presence of an evaluative property, given that she's properly shaped and situated, then she has generative emotional competence.

The problem is that the scope of success in this conditional is ambiguous. To illustrate, suppose that, as a dedicated canine behaviourist, I am extremely reliable when it comes to experiencing fear with respect to fearsome dogs. Moreover, I am also reliably successful at taking these specific experiences at face value and forming evaluative beliefs that particular dogs are fearsome on the basis of those experiences. However, suppose that I am *drastically* unreliable when it comes to experiencing appropriate fear in response to non-dog fearsome objects; my reliability when it comes to appropriate fear experiences is very strictly localised to instances of fearsome dogs. Moreover, we can stipulate further that I am also drastically unreliable when it comes to undergoing non-fear emotional experiences that are appropriately responsive to the presence of an evaluative property. I almost never experience shame in response to the shameful, joy in response to the joyful, admiration in response to the admirable, and so forth. Thus, we have a case in which I'm likely to succeed in experiencing fear which is appropriately responsive to the fearsomeness instantiated by particular dogs, but very unlikely to succeed in experiencing an appropriate emotional experience outside of this specific scenario.

Now, assuming that my proper shape and situation remains constant across these cases (i.e. my comparative unreliability with respect to the other emotional spheres cannot be explained in terms of my lacking proper shape and situation), there's an open question as to whether I count as possessing generative emotional competence. In response to this ambiguity, I take it that there are three ways in which the defender of EC might attempt to clarify the view. We can spell out these options as follows:

- (i) Global Emotional Competence: If S is likely to succeed in experiencing *any* emotion appropriately responsive to the presence of an evaluative property, given that she's properly shaped and situated, then she has generative emotional competence.
- (ii) Type-Local Emotional Competence: If S is likely to succeed in experiencing a given emotion-type (e.g. fear) appropriately responsive to the presence of the relevant evaluative property (e.g. fearsomeness), given that she's properly shaped and situated,

then she has generative emotional competence with respect to that particular emotion-type.

(iii) Scenario-Local Emotional Competence: If S is likely to succeed in experiencing a given emotion-type (e.g. fear) appropriately responsive to the presence of the relevant evaluative property (e.g. fearsomeness) within the context of a particular eliciting scenario (e.g. with respect to encountering fearsome dogs) given that she's properly shaped and situated, then she has generative emotional competence with respect to that particular emotion-type within the context of that particular eliciting scenario.

These options are presented in an ascending order of localisation. The Global option, insofar as the likely success calculation takes into account *all* of the subject's emotional experiences and the *entirety* of her emotional disposition, in effect renders generative emotional competence an all-or-nothing affair. Either the subject *is* likely to succeed in experiencing an appropriately responsive emotion, in which case the subject has generative emotional competence, or the subject *isn't* likely to succeed, in which case they don't. If they do have this global generative emotional competence, then, assuming the cooperation of the doxastic emotional competence, then *all* of the subject's emotional experiences will be capable of justifying their evaluative beliefs.

The localised options – Type-Local and Scenario-Local – deny this. For Type-Local, it's possible for a subject to possess generative emotional competences with respect to particular emotion-types but not others. If I'm likely to succeed in experiencing anger appropriately responsive to the offensive, but unlikely to succeed in experiencing envy appropriately responsive to the enviable, then my experiences of anger are capable of justifying my evaluative beliefs about offence, but my experiences of envy are incapable of justifying my beliefs about the enviable. For Scenario-Local, the calculations of likely success and the corresponding competence verdicts are localised even further. On this interpretive option, it's possible for a subject to possess generative emotional competence with respect to certain emotion-types elicited by *some* scenarios but not others. If I'm likely to succeed in experiencing anger appropriately responsive to offensive comments from my peers but unlikely to succeed in experiencing anger appropriately responsive to offensive and unjust political policies, then my comment-elicited experiences of anger are capable of justifying my evaluative beliefs about the offensiveness of those comments, but my policy-elicited

experiences of anger are incapable of justifying my beliefs about the offensiveness of those policies.

Hopefully it's clear, then, that EC's ambiguity in the scope of competences matters. How the defender of the view elects to clarify how we ought to make the calculation of likely success will change the ultimate justificatory result of certain emotional experiences. Taking the illustrative case above, Global will determine that my experiences of fear in response to fearsome dogs, despite being reliably responsive to the presence of fearsomeness, will *not* be capable of justifying my corresponding evaluative belief given that I'm not *overall* likely to succeed in experiencing appropriately responsive emotions, and so fail to possess generative emotional competence. Similarly, Type-Local will determine that my experience of fear with respect to fearsome dogs will *not* be capable of justifying my corresponding evaluative belief given that I'm not likely to succeed in experiencing appropriately responsive fear *in general*, and so fail to possess generative emotional competence with respect to the emotion-type of fear. Only Scenario-Local generates the result that I possess generative emotional competence with respect to experiencing fear in response to fearsome dogs, such that these fear experiences are capable of justifying the relevant evaluative beliefs. Given this discrepancy across the different interpretive options, then, EC cannot be a fully informative view until this ambiguity is clarified.

So, the question arises: how do we go about clarifying the view? Immediately, I think that we can dismiss the Global option. One significant problem for construing emotional competence as a global, all-or-nothing affair is that endorsing such a view would make EC insufficiently nuanced to be able to offer a compelling account of emotional justification. There is something counterintuitive about hinging the justificatory ability of, say, one's anger experiences on whether one's experiences of shame and fear and joy are likely to succeed in being responsive to the presence of the relevant evaluative properties. Indeed, it seems very *prima facie* plausible that individuals have varying levels of reliability across different types of emotions in different circumstances.<sup>99</sup> If the likelihood of success is something that we

---

<sup>99</sup> One possible explanation for this variety of reliability across emotion-types may be that emotional experiences, as we saw in Chapter 4 (§4.3.1), are cognitively penetrable, i.e. they're non-trivially influenced by our background cognitive states and desires. If some emotion-types are cognitively penetrable to a greater degree than others (e.g. perhaps social emotion-types such as shame and envy are susceptible to a higher rate of cognitive penetration than basic emotion-types, such as fear and disgust), then this might explain the varying degrees of reliability across emotion-types. Of course,

care about when it comes to determining whether an emotional experience has justificatory ability, then, surely, it seems more natural to make more specific epistemic judgments about whether an individual's emotional experiences are capable of justifying evaluative belief according to this variety of reliability. Tarring all of one's emotional experiences with the same brush, i.e. having an all-or-nothing determination as to whether one possesses generative emotional competence, and thereby whether one's emotional experiences are capable of justifying her evaluative beliefs, doesn't take into account the complexity of our emotional dispositions.<sup>100</sup> For this reason, then, I propose that we dismiss the Global interpretation of the trigger-manifestation conditional for generative emotional competence, and turn our attention to the localised variants as the more plausible answers.

Accordingly, the remainder of this chapter will endeavour to show two things. First, I will argue that the defender of EC ought to advocate on behalf of the Scenario-Local option: that generative emotional competence is localised with respect to particular emotion-types experienced in response to particular eliciting scenarios, such that the relevant trigger-manifestation conditional concerns whether the subject is likely to succeed in experiencing a given emotion-type appropriately responsive to the presence of the relevant evaluative property within the context of a particular eliciting scenario. As will become clear, this is because epistemically valuable generative emotional competences develop via processes of emotional learning, and these learning processes are localised with respect to particular eliciting scenarios. Second, and as a result of this conclusion, I will argue that this motivates a refined version of EC grounded in the notion of emotional learning. I will then attempt to defend this crystallised account as the most plausible theory of emotional justification considered thus far.

---

recall from §4.3.1 that an appeal to the cognitive penetrability of a particular emotion-type *by itself* is insufficient to deliver the result that that emotion-type will be altogether unreliable.

<sup>100</sup> Note also that the Global option fails to mirror our intuitive judgments elsewhere. Suppose, for instance, that we're assessing a particular subject for perceptual competence. Now, this subject has exceptional visual acuity, and is thereby extremely likely to succeed in having an accurate visual experience of her surroundings. However, suppose further that the subject's auditory and gustatory acuities are comparatively very low, such that she is very *unlikely* to succeed in undergoing accurate auditory and gustatory experiences. If a verdict of perceptual competence required that the subject was *overall* likely to succeed in having an accurate perceptual experience (i.e. encompassing all sensory faculties), then our subject would presumably fail to count as possessing perceptual competence given her drastic unlikelihood of success in the non-visual sensory domains. As such, an analogous view to Global EC concerning the epistemology of perception would determine that the subject's visual experiences were unable to justify the relevant beliefs given her overall lack of perceptual competence. This seems like the wrong result.



### 6.5. Solution: Emotional Competences Localised to Eliciting Scenarios

At this point in the dialectic, we're faced with the following question: is the relevant level of competence-localisation relative to the particular emotion-type experienced by the subject, or is it relative to the particular emotion-type experienced by the subject in the context of a particular eliciting scenario? In other words, is it only possible to possess generative emotional competence with respect to, say, fear experiences, or can one possess generative emotional competence localised further to fear experienced within the context of a particular eliciting scenario? It is my view that closer inspection into the development of generative emotional competence illuminates a plausible and motivated answer to this question. More specifically, I propose that the scope of generative emotional competence ought to be localised with respect to particular eliciting scenarios given the way in which these competences develop in the first place.

Now, interestingly, we've already come across a contender for a plausible theory as to how individuals come to possess at least *some* generative emotional competences. Recall from Chapter 2 (§2.4.1) that a compelling explanation as to how we come to learn which emotional experiences are appropriate within the context of particular situations is de Sousa's sketch of emotional learning. As a quick recap, de Sousa's view is that individuals come to learn which emotions are appropriate by gradual acquaintance with particular formative emotional contexts, i.e. "paradigm scenarios" (1987, p. 434). These scenarios are experienced in infancy and involve the subject experiencing instinctive responses towards particular objects and events in a given situation. The subject then comes to relate their experience to the particular features of the situation-type and thereby comes to grasp emotion concepts insofar as they associate their emotional experience with the eliciting scenario.

What makes this relevant for our current purposes is that this view strikes me as capturing *at least* something close to the development of generative emotional competence insofar as it paints a picture in which individuals, through time and association, come to learn that certain emotions are appropriately associated with specific types of eliciting scenarios (i.e. scenarios which instantiate particular sets of non-evaluative properties). This type of emotional learning strikes me as at least indicative of coming to learn that certain eliciting scenarios (i.e. certain objects and states of affairs) instantiate particular evaluative properties, and, hence,

becoming reliably successful at experiencing emotions appropriately responsive to the presence of those evaluative properties.

Crucially, however, de Sousa's account is limited to a focus on the development of emotional learning throughout childhood and infancy. I take it that the development of generative emotional competence, on the other hand, is not restricted to one's early years. It seems very plausible that we can come to emotionally learn that certain contexts and objects instantiate particular evaluative properties (such that we come to learn which emotions are appropriate within the context of particular scenarios) throughout the course of our lifetimes. Take the canine behaviourist who, through prolonged study of and exposure to different canine behaviours, comes to learn which dogs instantiate the property of fearsomeness, such that she reliably experiences fear in response to the fearsome-making features of dogs (e.g. particular aggressive behavioural indicators). This is a type of emotional expertise which arises not merely from early childhood experiences, but from more complex associations forged later on in life.

Inspired by an expanded account of de Sousa's theory, then, we can hypothesise the following: generative emotional competence (i.e. the ability to be reliably successful in experiencing an emotion appropriately responsive to the presence of evaluative properties) develops via gradual acquaintance and habituation with particular eliciting scenarios, such that, over time, the subject comes to emotionally learn that certain objects and states of affairs instantiate particular evaluative properties.<sup>101</sup> This, I take it, provides us with the resources to provide a motivated answer to the question that we started with at the beginning of this section. Generative emotional competences are localised to particular emotions experienced in response to particular eliciting scenarios *because* these competences develop within the contexts of these particular scenarios through mechanisms of emotional learning.<sup>102</sup>

---

<sup>101</sup> Interestingly, this proposal strikes me as cohering particularly well with the Aristotelian theory concerning the development of virtue considered in Chapter 4 (§4.3.3). Recall that, according to Aristotelian virtue ethics, virtues are character traits which crucially involve specific emotional dispositions to undergo appropriate emotional experiences within particular contexts, and these virtues develop via prolonged practice and habituation. Hence, this looks like a theory which proposes that we come to develop emotional dispositions that are reliably responsive to the presence of evaluative properties via prolonged practice with and exposure to specific eliciting scenarios. Plausibly, the development of these emotional dispositions corresponds to the develop of generative emotional competences.

<sup>102</sup> This rough proposal will be refined and clarified in the following sections.

What does this mean for our theory of emotional justification? I take it that we can paint a picture of emotional justification corresponding to this construal of generative emotional competences as follows. At a time  $t^1$ , a subject S undergoes an emotional experience which attributes a particular evaluative property E to an object (or state of affairs) O. Now, assuming the cooperation of doxastic emotional competence, S's emotional experience is capable of immediately and defeasibly justifying S's corresponding evaluative belief insofar as S has undergone processes of emotional learning, such that she has developed a learned generative emotional competence to reliably undergo emotional experiences which attribute E to O, and, crucially, S's emotional experience arises from this learned competence. Putting this into practice, our canine behaviourist's experience of fear with respect to a particular dog is capable of immediately and defeasibly justifying her evaluative belief that the dog is fearsome if and only if: she has emotionally learned that dogs with these particular features instantiate the evaluative property of fearsomeness; her experience of fear arises via this learned generative emotional competence; and she forms her evaluative belief on the basis of this emotional experience.

So, combining EC with the idea that it is through the process of emotional learning that one develops emotional competences, I propose the following view of emotional justification:

Learned Emotional Competence View (LEC): S's emotional experience which attributes evaluative property E to object O is capable of immediately and defeasibly justifying her belief that O is E if and only if (i) S has the emotional experience as a result of her having emotionally learned that O is E, and (ii) S forms the belief that O is E on the basis of that emotion.

The first condition in this formulation corresponds to the subject's possession of generative emotional competence, while the second corresponds to her possession of doxastic emotional competence. That is, the only substantive difference between LEC and EC is that the former spells out generative and doxastic emotional competence in terms of the subject's learned emotional abilities. The former, then, is a refined version of the latter. So, in a nutshell, *only* emotional experiences which arise as a result of the subject's learned capacity to emotionally identify an object as instantiating the given evaluative property are capable of justifying the relevant evaluative belief, and it's the particular learned capacity which determines the subject's epistemically relevant emotional competence. Emotional experiences which are

generated instead by false beliefs, poor reasoning, etc. are *not* capable of performing an immediate justificatory role.

This is the theory of emotional justification that this dissertation will defend as the most promising account considered thus far. In what follows, I will expand upon this view by further clarifying and elucidating the notion and mechanisms of emotional learning in §6.5.1, identifying a unique selling point of LEC in §6.5.2, and summarising the strengths of the view in §6.6.

### 6.5.1 The Processes of Emotional Learning

Let me begin by providing a more thorough characterisation of the mechanisms and processes of emotional learning. In the first instance, it may be instructive to elucidate the notion of emotional learning by analogy to perceptual learning, i.e. the process through which long-term changes are made to perception as a result of practice and experience.

Paradigmatic examples of perceptual learning typically include the development of particular expertise. For example, take a wine aficionado's ability to perceptually discern the quality of wine according to taste. Here, through prolonged exposure to the qualities of different wines, the aficionado comes to be perceptually sensitive to the subtle low-level properties which determine the quality of the wine. Importantly, note that perceptual learning differs from perceptual *development* insofar as the process of perceptual learning can occur at any point throughout one's life; it's not restricted to development in one's early years. *Emotional* learning, on the other hand, plausibly happens via a similar mechanism. Through repeated experience with and exposure to particular situations, a subject comes to emotionally learn (via association with their past experience) which objects instantiate particular evaluative properties, i.e. which emotional responses are fitting with respect to the given object.

What's essential for our purposes is the epistemic significance of these learning mechanisms. Many take skills and expertise brought about by the process of perceptual learning to be epistemically significant when it comes to the justificatory status of the relevant beliefs. Take Brogaard and Gatzia's (2018) claim that, in virtue of giving rise to particular perceptual expert skill-sets, "perceptual learning can thus turn our appearances into strong immediate justifiers of the beliefs we form as a result of said learning" (p. 554). Or, on a slightly different tack, recall Markie's (2013) view that an agent's seeming experience that something

is Q is capable of bearing justificatory power if and only if that experience “is had in the exercise of the subject’s knowledge of how to perceptually identify something as being Q” (p. 262).<sup>103</sup> While Markie doesn’t specifically mention the process of perceptual learning, it’s plausible that one’s possession of a knowledge-how capacity to perceptually identify an object as being Q either amounts to (or is at least substantively dependent upon) having perceptually learned that objects that look like *that* are Q.<sup>104</sup>

From these views, we can see that drawing a connection between one’s expertise brought about by the mechanisms of perceptual learning and the justificatory status of one’s relevant belief is not unfounded. To motivate this, suppose that you and I spend a pleasant afternoon wine-tasting. You, a competent and experienced winetaster, have learned to perceptually identify merlots by their taste properties, whereas I have no such perceptual expertise. At a time  $t^1$ , you and I both have a perceptual experience of the same wine. In virtue of your learned capacity, you are able to perceptually identify the wine as being a merlot, i.e. long-term changes have been made to your perceptual system through practice and experience, such that you are able to attend to the merlot-making features of the wine. This is not the case for me. At best, I can have a perceptual experience with the contents ‘there is a wine that tastes *like that*’, and guess that wines that taste *like that* are merlots, but I cannot have the perceptual experience with the contents ‘that is a merlot’, whereas, plausibly, you can. Perceptual learning, then, is feasibly what explains the intuitive verdict that your perceptual experience is capable of immediately justifying your belief that the wine is a merlot, while mine does not.

LEC tells an analogous epistemological story with respect to emotional learning and the justificatory power of emotional experience.<sup>105</sup> That is, just as the expert wine-taster enjoys

---

<sup>103</sup> Remember from Chapter 3 (§3.3.2) that, for Markie, a knowledge-how capacity to identify objects as being Q amounts (at least partially) to the possession of the relevant background information, i.e. justified beliefs that would enable one to identify an object as being Q. This raises difficulties for understanding the justification conferred as being *immediate*, but this is not a problem for my view given that LEC does not strictly mirror Markie’s account.

<sup>104</sup> Both Brogaard and Gatzia and Markie endorse forms of dogmatism about perceptual justification. Given the failures of dogmatism identified in Chapters 2 and 3, I do not ally myself with dogmatism here; I only use these views to highlight the fact that learning mechanisms have been considered crucial to understanding immediate perceptual justification in the literature.

<sup>105</sup> Note, however, that I do not hinge the plausibility of LEC on there being a strict analogy with perceptual learning. The analogy is intended only to illuminate what I take to be the central features and epistemic significance of emotional learning processes.

immediate justification for their belief that the wine is a merlot in virtue of their having perceptually learned that wines that are perceptually *like that* are merlots, an individual who experiences fear in response to an approaching snake enjoys immediate justification for their evaluative belief that the snake is fearsome in virtue of their having emotionally learned that objects *like that*, i.e. objects with those fearsome-making non-evaluative properties, are fearsome.

Now, as a matter of contrast with perceptual learning, the notion of emotional learning has not, as of yet, been fully explored with respect to the ability of emotional experience to immediately and defeasibly justify evaluative belief. In spite of this, however, I take it that there *are* notes of indirect support for this idea in the contemporary literature. Echeverri (2019), for example, draws an analogy between perceptual and emotional recognition (or learning) in his discussion of how we can conceive of the emotions being immediately justified *themselves*, i.e. how emotional experience can be immediately justified by the contents of their cognitive bases. Put roughly, Echeverri argues that a subject's emotional experience is immediately justified if, through the cognitive base of the emotion, she is aware that a particular object instantiates a set of non-evaluative properties, and she has the emotional disposition that would cause her to seamlessly "move from the detection of" (p. 560) the non-evaluative properties to the relevant emotional experience. Crucially, the point of similarity between Echeverri's view and mine is that Echeverri argues that the transition between cognitive base awareness of non-evaluative properties to emotional awareness of the corresponding evaluative property is "phenomenally immediate", and brought about by processes of recognition and learning:

Expert birdwatchers have achieved high reliability in bird recognition. This high reliability is reflected in the phenomenology of perceptual recognition, i.e. in how expert birdwatchers experience the transition from the representation of low-level properties to the representation of high-level properties. Something similar occurs in our most basic emotional responses. Presumably, our basic emotional dispositions have developed through long evolutionary processes in relatively stable environments. Some of these emotional dispositions can even be shaped and modified through learning, therapy, or habit... Once we are endowed with the relevant dispositions, the emotional responses they ground are permitted from the first-person perspective. (p. 559)

On Echeverri's view, the epistemic relevance of emotional learning processes is that these processes plausibly forge emotional dispositions which allow us to transition from perceiving the non-evaluative properties instantiated by an object to experiencing an emotion which provides access to the relevant evaluative property in a phenomenally immediate way. Of course, given that my focus is on the question of how emotional experience can immediately justify evaluative belief, it is beyond the scope of this dissertation to provide any substantive analysis on Echeverri's view of how emotions *themselves* can be immediately justified. However, what is important for our purposes here is to note that emotional learning has been appealed to as part of the explanation as to how epistemic justification can attach to emotional experiences. Moreover, as we'll see later, I take it that we can also find supportive notes for the possibility of epistemically significant emotional learning processes in the feminist philosophy literature. §6.5.2 will be dedicated to elucidating this connection, and highlighting the theoretical advantages that LEC bears with respect to insight from this literature.

Importantly, however, LEC does not *just* enjoy support from related views within the philosophical literature. A major selling point of a view of emotional justification grounded in the possession of learned emotional competences is that coheres with a host of empirical support from related views and findings in developmental psychology and cognitive science. For example, consider the following from Pollak et al (2009). In pursuit of uncovering how the early social experiences of children influence their abilities to recognise the emotional states of others (as indicated by their facial expressions), it was found that children raised in physically abusive environments, i.e. environments marred by frequent episodes of hostility, were typically able to accurately identify expressions of anger much faster, and on the basis of significantly less visual information, than non-abused children. Pollak et al. summarise the findings as follows:

The present experiment examined children's construction of emotion representations from fragmentary or partial information as models generated naturalistic facial expressions. We found that children who have been exposed to unusually high levels of anger were able to accurately recognize anger early in the formation of the facial expression of anger, when fewer expressive cues (such as activation of facial musculature) were available. However, these children performed similarly to controls

when viewing other emotional expressions, suggesting that the abused children neither had a bias to select anger nor that the abused group generally performed or understood the task differently than did controls. (p. 243).

There are a couple of important things to note about these findings. First, Pollak et al. present the developed perceptual abilities of these children as an instance of perceptual learning. That is, through prolonged exposure to instances of expressed anger and hostility, the perceptual systems of children raised in abusive environments have become sensitised to recognising specific kinds of perceptual information, i.e. facial cues indicative of anger.<sup>106</sup> Second, note that this perceptual ability is strictly localised to recognising expressions of anger. The abusive environment did not prompt long-term changes in sensitivity to perceptual information conveyed by facial cues expressive of surprise, or joy, for instance.

Now, what makes these findings relevant for our purposes is that, plausibly, this developed perceptual ability is not *only* an instance of perceptual learning, but also at least partially comprises an instance of emotional learning. When stipulating why the perceptual systems of abused children develop this sensitivity towards specific visual information, Pollak et al. suggest the following:

The development of increased perceptual sensitivity for the fine-grained details of variation in affective expressions may provide a behavioral advantage for children living in threatening contexts, allowing earlier identification of salient emotions. Anger-related cues may become especially salient to physically abused children *because they are associated with imminent harm*. (emphasis added; p. 246).

It is this association with imminent harm that strikes me as suggestive of a developed emotional capacity to identify instances of fearsomeness. Forging an association between anger-related cues and harm seems to me to be a plausible instance of emotionally learning that a particular set of non-evaluative properties instantiated in a given eliciting scenario (i.e. visual properties instantiated by angry-looking facial expressions) give rise to the evaluative

---

<sup>106</sup> A similar verdict was delivered by a study conducted by Wilkowski and Robinson (2012). Here, roughly speaking, it was shown that physically aggressive individuals had an increased perceptual sensitivity to angry facial expressions compared to non-aggressive individuals. Exposure to aggressive environments, such as harsh parenting, was noted as a significant causal contributor towards physical aggression (see p. 542).



property of fearsomeness. So, given this development of particular competences to perceptually identify angry facial expressions, and the likely corresponding development of competences to emotionally identify fearsome expressive and behavioural traits, it strikes me as very plausible that these children's experiences of fear in response to particular facial cues are capable of justifying their evaluative beliefs concerning the fearsomeness of their situation and their aggressor. Crucially, again, the justification conferred by emotional experiences which arise via the relevant learning processes is immediate; it doesn't depend on any inferential connections to any existing justified beliefs, nor does it require the children to have an awareness of the relevant non-evaluative visual properties or their perceptual ability. What I take this research to show, then, is that there is compelling empirical evidence for the proposal that individuals can and often do develop localised competences with respect to undergoing accurate perceptual experiences, and the development of these competences is facilitated by processes of perceptual learning. Moreover, these perceptual learning processes are intricately bound up with the way in which the individual construes their evaluative environment. My view of emotional learning, insofar as it involves a subject becoming sensitised to the instantiation of particular evaluative properties in a given eliciting scenario (via becoming sensitised to the non-evaluative properties which give rise to the evaluative property in question), coheres well with these findings.

In summary, I take it that the notion and epistemic significance of emotional learning enjoys support both from close similarities and connections to processes of perceptual learning, and from closely related philosophical and empirical research. The process of emotional learning takes place over time and through repeated exposure to specific eliciting scenarios. The subject, through patterns of association and learning, comes to be emotionally sensitised to the presence of evaluative properties through becoming sensitised to the relevant sets of non-evaluative properties. Crucially, this process is not cognitively demanding; the subject need not be aware of this process taking place. As suggested by Pollak's findings, children can and do undergo processes of emotional learning, such that their emotional experiences which arise via these processes are capable of conferring immediate justification to their evaluative beliefs (provided these beliefs are formed on the basis of these emotional experiences).

### 6.5.2 Emotional Learning and the Insight from Oppression

Having motivated and clarified LEC, let me now highlight what I take to be a powerful selling point for the view. Namely, I argue that LEC is uniquely capable of explaining a plausible insight defended by many in the feminist philosophy literature, namely, that certain emotions experienced by individuals from oppressed socio-political groups are capable of bearing epistemic value in virtue of those individuals' social standing.<sup>107</sup> Let me begin by contextualising and clarifying this insight, before explaining how LEC manages to capture it.

The idea that certain emotions experienced by members of oppressed groups can serve as valuable epistemic resources is a common one in feminist epistemology. As a clear statement of such a view, consider the following influential remarks from Jaggar (1989):

Subordinated people have a kind of epistemological privilege in so far as they have easier access to this standpoint [of the subordinated] and therefore a better chance of ascertaining the possible beginnings of a society in which all could thrive. For this reason, I would claim that the emotional responses of subordinated people in general, and often of women in particular, are more likely to be appropriate than the emotional responses of the dominant class. That is, they are more likely to incorporate reliable appraisals of situations. (p. 168)

There are a couple of important things to note here. First, note the emphasis on the increased likelihood of emotional reliability within the emotional dispositions of oppressed persons. Plausibly, what Jaggar has in mind here is that emotions experienced by members of oppressed groups are especially likely to succeed in being appropriately responsive to particular evaluative properties. Of course, this does not entail infallibility. Moreover, note that the explanation for this is that subordinated or oppressed individuals have unique access

---

<sup>107</sup> Silva (2020) also discusses a version of this insight with respect to a theory of emotional justification. More specifically, Silva focuses on the epistemic value of outlaw emotions (i.e. emotional experiences which bear representational contents that conflict with dominant social values), and argues that accommodating for this epistemic value requires advocating on behalf of a reliabilist version of the justificatory thesis informed by standpoint theory. Silva's project is similar to mine, then, insofar as she takes likely success brought about by occupying a particular epistemological standpoint to be essential for the emotions' capacity to justify evaluative belief, but differs from mine in that she does not explain or motivate the significance of standpoint reliability in terms of emotional learning processes.

to a particular perspective, or standpoint, through which one can accurately apprehend features of her oppressive environment. It is *in virtue* of being oppressed that a marginalised individual has increased likelihood of reliability when it comes to emotionally apprehending relevant features of her environment.

To illustrate, consider a subject who experiences sexualising remarks at the hands of her colleagues. Their sexism and misogyny manifests in different ways, such as commenting on the way in which she dresses, making sexualising jokes about her body, and so forth. Further, suppose that, within her workplace, comments such as these are considered to be innocuous. The general social ethos is that remarks of this sort constitute good-natured horseplay, rather than instances of serious harassment and degradation. Despite this ethos, however, the subject experiences a degree of anger when she's subjected to these comments. Jaggar's view is that an emotional experience of this sort bears distinct epistemic value. More specifically, the subject's anger has a high likelihood of being successfully responsive to the evaluative property of 'offensiveness' instantiated by the remarks (i.e. a high likelihood of accurately representing the comments as being offensive) in virtue of her having unique epistemic access to the standpoint of victims of harassment.<sup>108</sup>

Importantly, it is not just Jaggar who maintains this view. Narayan (1988), as further notable example, argues that "the emotions play an important role in the knowledge that is part of the epistemic privilege of the oppressed" (p. 36), and, moreover, that these epistemically valuable emotions are brought about through the individual's "lived experience of oppression" (p. 39). It is a well-received view in feminist philosophy, then, that the emotional experiences of oppressed individuals are uniquely epistemically valuable insofar as they can reliably represent aspects of their evaluative situation; they can shed light on evaluative facts otherwise obscured by the oppressive climate.

In the interest of clarity, we can set out these two important claims as follows:

- (a) Certain emotions experienced by members of oppressed groups bear significant epistemic value.

---

<sup>108</sup> Later in this section, I clarify and expand upon this claim, and provide my own explanation as to why I think that it is in virtue of *being* oppressed that oppressed individuals have epistemically significant emotional experiences.

- (b) These emotions bear this epistemic value *because of* the emoter's social standing, i.e. their membership of the oppressed group.

Now, at this point, there is a further question to be addressed. Namely, we must determine what the 'epistemic value' possessed by emotions experienced by oppressed individuals actually amounts to. Here's one proposal. Borrowing argumentative resources from Brady (2013), it might be suggested that the distinct epistemic value of emotions experienced by members of oppressed groups is their capacity to prompt reflection and revision of evaluative belief through motivating the search for further reasons. Reconsider the case above. We might think that what's epistemically valuable about the subject's anger towards her colleagues and their sexualising comments is that the experience prompts her to closely inspect and potentially revise her societally-sustained evaluative belief that remarks of this sort constitute good-natured horseplay. That is, what's valuable about her anger is that it sets off a causal chain through which she begins to search for further evidence which bears on her anger; to look for further evidence in favour of the comments actually constituting an offence, rather than innocuous social banter. Emotions experienced by oppressed individuals throw a spanner in the works of the individual's societally-sustained belief system about their oppressive circumstances, so to speak, and in prompting the deconstruction of this belief system through motivating the search for further reasons, they have epistemic value.

Now, perhaps this is one of the ways in which emotions experienced by oppressed individuals are useful. Having said that, I don't think that this proposal is a plausible explanation of the distinct epistemic value of these emotional experiences as gestured towards above by Jaggar and Narayan. The central reason for this is that it's not clear that this proposal captures the idea that emotions experienced by oppressed persons are distinctly *epistemically* valuable. The intuition gleaned from these insights from feminist philosophy, I take it, is that emotions act as unique sources of epistemic evidence or information that are essential for the formation of true evaluative beliefs about objects and states of affairs in the oppressed subject's surroundings. Indeed, recall Jaggar's claim that these emotions "are more likely to be appropriate than the emotional responses of the dominant class". This suggests that these emotions' increased likelihood of reliability is part and parcel of their epistemic value. The idea that their epistemic value is rooted in their causal ability to initiate processes of belief-revision doesn't quite seem to capture this insight.

So, if this epistemic value can't be fully explained in terms of prompting an epistemically valuable causal chain of belief-revision, how *can* we identify this epistemic value? I think that the limitation of the above proposal illuminates a more plausible answer. Namely, that the epistemic value of emotions experienced by oppressed persons is, in fact, best understood in terms of their possessing immediate justificatory power. The proposal that the distinct epistemic value of these emotions amounts to their capacity to immediately and defeasibly justify evaluative belief fits much better with Jaggar's remarks, given the relevance of reliability and accuracy to the question as to whether an experience can justify belief. Thus, from the above, we have the following three claims:

- (a) Certain emotions experienced by members of oppressed groups bear significant epistemic value.
- (b) These emotions bear this epistemic value *because of* the emoter's social standing, i.e. their membership of the oppressed group.
- (c) This 'epistemic value' possessed by emotions experienced by members of oppressed groups amounts to immediate justificatory power.

Let us refer to the conjunction of these three claims as 'the insight from oppression'. Given the philosophical clout that this insight carries in feminist epistemology, and given its plausibility, it strikes me that a plausible theory of emotional justification ought to be able to accommodate for this insight. That is, a plausible version of the justificatory thesis of emotion will be able to explain why certain emotions experienced by oppressed individuals are capable of immediately and defeasibly justifying their evaluative beliefs.

Now, interestingly, the views of emotional justification considered earlier in this dissertation do not provide a particularly promising explanation of this epistemic value. Take emotional reliabilism (ER). ER can only determine that these experiences are capable of justifying the relevant evaluative beliefs if and only if forming true evaluative beliefs on the basis of their emotional experiences is a reliable process of belief-formation. The problem for ER with respect to the case of oppressed individuals, however, is that they plausibly do *not* reliably form true evaluative beliefs on the basis of their emotional experiences.<sup>109</sup> We've seen that an

---

<sup>109</sup> This effect of an oppressive climate on subordinated individuals strikes me as a plausible instance of (or, at least, a phenomenon closely related to) epistemic injustice. See Fricker (2007).

oppressed individual's social standing can render them more likely to undergo specific emotional experiences that are accurately responsive to the presence of evaluative properties in specific scenarios, but it strikes me as equally plausible that an oppressive climate can distort a subject's emotional disposition and render them overall more likely to experience emotional mistakes in alternative scenarios.<sup>110</sup>

As an illustrative example, take a member of a racial minority group whose name is continually and habitually mispronounced by those in his community. These occasions happen frequently across his life, marking not only his interactions and communications in the social domain, but also within the sphere of his professional life. He thereby comes to identify these mispronunciations not just as harmless or accidental slips of the tongue, but rather as token instances of a general type, i.e. systematic devaluing of his culture and heritage. As such, he reliably experiences some degree of apt indignation upon hearing his name mispronounced while others around him, from non-minority racial groups, fail to experience such indignation. However, suppose further that, because of the oppressive climate, the agent is also particularly susceptible to inaccurately experiencing guilt when correcting or confronting the perpetrators of the microaggressions, or inaccurately experiencing admiration towards members of non-minority racial groups who correctly pronounce his name. It seems very plausible that, overall, an oppressive climate renders the emotional belief-forming process of this individual unreliable, such that ER would be unable to bestow justificatory power to emotions experienced by similar individuals.

Likewise, I take it that both Global and Type-Local instantiations of EC will fall foul of a similar worry. Specifically, the oppressed subject's comparative unlikelihood of success with respect to many other instances of emotional experience (e.g. their guilt and admiration experiences) will unfavourably tip the scale, such that the subject will likely fail to meet the requirement specified in the trigger-manifestation conditional (i.e. the Global requirement that the subject is overall likely to succeed in experiencing an appropriately responsive

---

<sup>110</sup> This, I take it, provides a motivated answer to a generality-style challenge facing ER. Specifically, the idea is that ER ought to construe the relevant belief-forming process narrowly given the discrepancy of reliability one can have across different types of emotional experiences and across different eliciting scenarios. Of course, even ER qualified in this way will still face troubling clairvoyance-type challenges elsewhere, but I take it the reliabilist approach ought not be rejected on the grounds of being unable to answer generality-style challenges, as I argued in Chapter 4 (§4.2.2).

emotion or the Type-Local requirement that the subject is likely to succeed in experiencing a given emotion-type that is appropriately responsive to the presence of evaluative properties).

Emotional dogmatism (ED), on the other hand, likely will be able to bestow justificatory power to emotions experienced by oppressed individuals (assuming that these emotional experiences meet the phenomenological requirements of the view), but it doesn't seem to be able to *explain* the epistemic significance of these emotional experiences in a particularly compelling way. Recall that all that matters for the dogmatist, epistemically speaking, is that the emotional experience in question has the relevant phenomenological character. Moreover, recall from Chapter 2 (§2.2.1) and Chapter 3 (§3.2.1) that many emotional experiences which we don't take to be capable of bearing justificatory power (e.g. the biased interviewer's emotional experience) meet the phenomenological requirements of ED. So, although ED may be able to accommodate for the epistemic significance of emotions experienced by members of oppressed groups, it can't provide a plausible or discerning explanation of this significance given that emotions like those of the biased interviewer will *also* bear the same significance on the account.

In contrast to the shortcomings of the above views, LEC bears the resources to both explain and further illuminate the insight from oppression. My view is that LEC accommodates for this insight by neatly fitting into the explanation as to *why* group membership comes alongside the development of epistemically significant emotional experiences. On this note, consider the following remarks from Narayan (1988) concerning the difference in epistemic perspective between 'insiders' of oppressed groups (i.e. oppressed individuals), and the privileged 'outsiders' (i.e. members of dominant social groups):

An outsider who has not experienced oppression first-hand and has learned about it second-hand, is more likely to understand the general and commonplace ways in which the oppression is manifested. For instance, if a professor uses openly racist or sexist examples..., sympathetic white male students may be able to spot his attitudes quite as well as the victims of the attitude. But if his attitudes are expressed more covertly, through dismissing their queries, not taking their contributions seriously, undervaluing their work, lack of cordiality, etc. outsiders may fail to see what is going on. *An insider who is sensitised to such prejudiced attitudes will often pick up on cues ranging from facial expressions to body language that an outsider may simply fail to*

*spot and will often also be alerted by her own feelings of unease about the person or situation.* (p. 39, my italics)

I believe that Narayan's comments here – particularly the italicised passage – succinctly capture what it is about oppressed group membership that enables the development of epistemically significant emotional experiences. Oppressed individuals, through time, become *sensitised* to the presence of oppressive attitudes and features of their social situation. That is, given the frequency with which oppressed individuals encounter oppressive attitudes, and given the significant negative impact that these attitudes have for such individuals, oppressed persons come to be sensitised to the presence of these attitudes. What this sensitisation amounts to, I propose, is a developed ability to apprehend subtle non-evaluative properties in their environment which give rise to particular evaluative properties, and *this* is what causes their emotional experiences to be reliable in being appropriately responsive to the presence of evaluative properties. This, of course, is an instance of my characterisation of emotional learning.

To illustrate, reconsider the harassment case. The reason that the subject's experience of anger in response to her harassers bears epistemic value (i.e. immediate justificatory power), is that, in virtue of her social-political status as a woman, she has, through experience, come to be sensitised to the ways in which sexualising comments *do* constitute an offence. That is, through time and experience, she comes to learn that these comments lead to her being devalued as a colleague, to her professional contributions being minimised, to her being humiliated in front of her peers, and so forth. Assuming, then, that actions and states of affairs which instantiate these particular non-evaluative properties consequently instantiate the evaluative property of 'offensiveness', she emotionally *learns* that sexualising comments instantiate the evaluative property of 'offensiveness' *through* becoming sensitised to the non-evaluative features of the comments that cause them to instantiate that evaluative property. Crucially, this process of sensitisation and learning is tied to her group membership as a woman. Outsiders to this group will be less likely to undergo these sensitisation and learning processes simply because they do not have the same sustained long-term experience with sexualising comments and their consequences. As a result of this, outsiders will likely fail to spot the subtle non-evaluative properties which give rise to the, say, offensiveness of sexualising comments.



At this point, let me be clear that I do not intend for group membership to be necessary *nor* sufficient for enabling the processes of sensitisation and learning. Taking the case of sexual harassment, there will, of course, be individuals who identify as women that do not undergo these processes of learning with respect to properties instantiated by instances of sexual harassment. Likewise, there may be individuals who do not identify as women that happen to undergo these processes themselves. My proposal is only that occupying a particular social standing aids in the facilitation of these processes, or puts one in the position such that she is more likely to undergo these processes of sensitisation and learning, such that her emotional experiences will be more likely to reliably track the presence of particular evaluative properties. Indeed, I think that the Pollak case discussed in §6.5.1 illuminates this. Recall that, here, we saw that children raised in a hostile environment were typically better sensitised to visual information conveyed by anger-expressive facial cues. The proposed explanation for this was that, for children raised in these environments, these anger-expressive facial cues are associated with imminent harm, such that early detection of these facial cues could be adaptive insofar as it may facilitate an early response to this threat. Put roughly, early detection of these facial cues *matters* for children in abusive environments in a way that it doesn't tend to for non-abused children. Similarly, I contend, detection of sexual harassment and its offensiveness *matters* for women (as typical victims of such harassment) in a way that it doesn't *tend to* for non-women, or in the same way that detection of racial microaggressions *matters* for individuals from minority racial groups in a way that it doesn't *tend to* for white individuals. Plausibly, then, the differential in this dimension of significance is what makes occupying a particular social standing capable of increasing or decreasing a subject's likelihood of undergoing the relevant sensitisation and learning processes.

Putting all of this together, then, here's the proposal: members of oppressed groups develop reliably successful emotional dispositions with respect to particular evaluative properties instantiated by particular states of affairs (e.g. the offensiveness of sexualising comments) because, in virtue of their particular social standing, they are well-placed to undergo the processes of *emotional learning* with respect to those states of affairs and the evaluative properties they instantiate. Hence, LEC can both accommodate for and illuminate the insight from feminist philosophy that certain emotions experienced by members of oppressed groups bear epistemic value.

Finally, note that LEC can also explain why it is not the case that *all* of an oppressed subject's emotional experiences are capable of bearing immediate justificatory power. Reconsider the individual who, as a result of being habitually subject to racial microaggressions, is particularly susceptible to inaccurately experiencing guilt when correcting or confronting the perpetrators of the microaggressions, or inaccurately experiencing admiration towards members of non-minority racial groups who correctly pronounce his name. LEC can explain why these guilt and admiration experiences cannot justify his corresponding evaluative beliefs, i.e. because they are *not* produced by the process of emotional learning, but instead upon internalised false beliefs about permissibility and praiseworthiness respectively in the domain of racial relations. That is, the subject has not undergone a long-term process through which he has become sensitive to the non-evaluative properties of, say, a person's pronouncing his name correctly which give rise to the evaluative property 'admirable'. Rather, his false background beliefs about the admirability of individuals who pronounce his name correctly have been caused and manipulated by the oppressive climate around him. Because of this, LEC does not bestow immediate justificatory power to these emotional experiences. LEC, then, seems to be particularly well-equipped for generating plausible and compelling results with respect to complex oppression cases.

## **6.6 The Virtues of The Learned Emotional Competence View**

Now, having motivated LEC by spelling out the process of emotional learning in more detail, and contextualising it with respect to similar ideas in the literature, let me end this chapter by summarising why LEC is, in my view, the most plausible account as to how emotional experience can immediately and defeasibly justify evaluative belief.

We've seen already that LEC provides us with a developed and compelling explanation as to how emotions experienced by members of oppressed groups are capable of immediately justifying their evaluative beliefs. Moreover, we've seen that LEC also bears the resources to be able to explain why not *all* of their emotional experiences are capable of bearing justificatory power. That is, LEC can both explain the epistemic significance of learned emotional experiences and accommodate for the general destructive effects of an oppressive climate on the reliability of one's emotional experiences. Importantly, though, it's not *just* with respect to oppression cases that LEC prevails. The view also meets the desiderata for a plausible theory of emotional justification that I identified back in §6.1. Recall that the first

desideratum was that our view excludes obvious over-generalisation cases from enjoying the epistemic yield. Taking the biased interviewer as an example, LEC is clear that the emotional experience of suspicion that arises from sexist bias is not capable of bearing justificatory power. Bias is not an instance of emotional learning – the biased subject does not, through time, come to be sensitive to the fact that ‘being a woman’ gives rise to the evaluative property of ‘duplicitousness’. Rather, the emotion is caused by a network of false beliefs and cognitive biases against women.

The second desideratum was that, on the other side of the coin, our view does not veer too far in the direction of restraint by excluding emotional experiences that we would ordinarily consider to be capable of justifying belief on spurious grounds. I think that LEC meets this desideratum. Certainly, LEC will end up being more restrictive than, say, emotional dogmatism. Given that emotional learning is a long-term process, and the variety of ways in which we can, and often do, form emotional experiences on the basis of *other* mental states and processes, such as bias, false beliefs, and so forth, there will likely be many emotional experiences that end up being ruled out of the account. Importantly, however, this exclusion does not strike me as spurious. Indeed, ruling out emotional experiences like this from bearing justificatory power strikes me as precisely the correct result, as stressed in the discussion of the first desideratum. Moreover, as we’ve seen above, plenty of emotional experiences that we *do* ordinarily consider capable of bearing justificatory power will be ruled into this account. Basic emotional experiences, such as fear of a snarling dog and disgust towards spoiled milk, will be capable of justifying belief given their basis in emotional learning processes. Plausibly, complex emotional experiences, such as aesthetic appreciation of an artwork or amusement towards a particular situation,<sup>111</sup> will also be capable of bearing justificatory power. It seems entirely possible that one can emotionally learn, say, that works of art which instantiate particular non-evaluative properties are aesthetically admirable, *even if* the subject cannot verbalise or elucidate what these non-evaluative properties are. Emotional learning does not necessitate any higher-order ability to explicitly and consciously identify or articulate the non-evaluative properties which give rise to the evaluative property in question. Thus, because both our simple and complex emotional

---

<sup>111</sup> Recall that these emotional experiences were objectionably excluded from an instantiation of Restricted Emotional Dogmatism (i.e. the non-evaluative property reading) in §3.2.2.

experiences can enjoy immediate justificatory power on this account, and because any exclusions are on epistemically reasonable grounds, LEC meets the second desideratum.

At this point, one might object that there is a class of emotional experiences which LEC cannot accommodate for, namely, emotional experiences which seem to be ‘hard-wired’ in our psyches. Take a fear of snakes, for instance. It might be argued that such a fear is not plausibly ‘learned’ through apprehending and subsequent sensitisation to the fearsome-making features of snakes, but, rather, fear of snakes is an evolutionary mechanism embedded within us. Indeed, one might argue that this hard-wired class of emotional experiences extends to the realm of disgust towards contaminated objects, too. If these emotional experiences do not arise via processes of emotional learning, then LEC will determine that they cannot be capable of bearing justificatory power. However, one might object that this seems objectionably restrictive. If this is the case, then LEC fails to meet the second desideratum.

Let us focus on the fear of snakes as the illustrative case. The central thing to say in response to this point is that such a response requires assuming that a fear of snakes is *entirely* hard-wired, such that there is *no* input from learned processes and associations. If it is the case that a fear of snakes is only facilitated by a hard-wired predisposition, but ultimately develops via processes of emotional learning, then LEC can determine such fear experiences as capable of bearing justificatory power. This is a complex empirical matter, of course, but note that there is notable support for the claim that learning processes are at least necessary for the development and manifestation of such fear experiences. Take Cook and Mineka’s (1990) assertion that “[a]lthough considerable controversy has existed over the extent to which this fear of snakes is “innate” as opposed to being based on learning, current evidence suggests that learning is necessary for the fear to manifest itself” (p. 373), or the following from Stevenson et al. (2014) that “there is a disposition to learn associations between particular animals, especially snakes, and the emotion of fear” (p. 326). Both of these remarks seem to suggest that, at most, organisms like us are evolutionarily predisposed to easily *learn* to fear snakes. If this is true, then this is consistent with LEC bestowing justificatory power to such fear experiences. All that needs to obtain for LEC to include these experiences is that they arise via processes of emotional learning, it does not require that there is *no* biological or evolutionary input into the facilitation of these processes. As such, I take it that, in lieu of compelling evidence that learning processes are not at *all* necessary for the development and

manifestation of emotional experiences such as a fear of snakes, LEC is not rendered unable to satisfy the second desideratum by an appeal to ‘hard-wired’ emotional reactions and experiences.

Moving on, recall that the third desideratum for a plausible theory of emotional justification is that it can explain the necessity but insufficiency of emotional reliability. This desideratum is straightforwardly met by LEC. It can account for the necessity of *some* degree of emotional reliability, given the necessity of likely success in experiencing an emotion appropriately responsive to the presence of an evaluative property. However, it also explains why reliability is not *enough*. Given that it is rooted in a virtue epistemological framework, it also necessitates that the reliability is a result of the subject having emotionally learned that a particular object instantiates the relevant evaluative property, i.e. that reliability is the result of a learned emotional competence to identify particular evaluative properties. As a result of this, LEC generates the right result in the clairvoyance-analogous caffeine case. The reason that my caffeine-caused evaluative beliefs are *not* justified, despite the reliability of my belief-forming process, is that my emotional experiences of anxiety do not arise as a result of long-term changes to the sensitivity of my emotional system. I don’t experience anxiety as a result of having emotionally learned that events of my working day which instantiate particular non-evaluative properties are anxiety-meriting, rather, I experience that anxiety as a result of having ingested too much caffeine.

For these reasons, then, I take it that LEC is a strong and compelling theory as to how emotional experience can immediately and defeasibly justify evaluative belief. It fares better than dogmatism insofar as it does not over-generalise to problematic cases, nor does it exclude epistemically viable emotional experiences on spurious grounds. It fares better than reliabilism insofar as it can both explain the necessity but insufficiency of emotional reliability for justification and generate the right result in the caffeine case. Finally, it fares better than unqualified EC insofar as its supplementary learning-based framework allows the account to evade the ambiguity facing the unqualified view, and to thereby provide a more informative and instructive epistemic verdict.

## 6.7 Concluding Remarks

The goal of this final chapter has been to refine and defend an agent-based view of emotional justification grounded in the possession of learned emotional competences. In §6.1, I began the chapter by taking stock of our discussion so far, and identifying three central desiderata that a plausible account of emotional justification must be able to satisfy. §6.2 then introduced agent-based views, and considered two distinct virtue epistemological approaches to epistemic justification as potential models for an agent-based view of emotional justification. The first of which was virtue responsibilism (§6.2.1), which placed epistemic focus on the subject's possession of intellectually virtuous character traits. I argued that virtue responsibilism could not provide a plausible model for an account of emotional justification given the incompatibility of the two epistemological projects. In light of this, I then suggested that Sosa's competence-based virtue reliabilist view (§6.2.2) provided us with the best starting point for an analogous view of emotional justification.

§6.3 then presented Carter's Emotional Competence view (EC), according to which S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if (i) that experience is formed aptly (i.e. manifests her generative emotional competence), and (ii) S forms *e* via the exercise of her doxastic emotional competence. I then argued that, while EC is a promising view insofar as it seems able to capture the desiderata identified for a plausible account of emotional justification, it must be clarified in order to be sharpened into the most compelling theory it can be. To this end, §6.4 identified an ambiguity in the success requirement for generative emotional competences. Specifically, I argued that EC was silent as to whether generative emotional competence required likely success across *all* emotional experiences (Global), likely success across all emotional experiences of a given emotion-type (Type-Local), or all emotional experiences of a given emotion-type experienced in response to a particular eliciting scenario (Scenario-Local). I argued that the Global option ought to be rejected given the complexities of our emotional dispositions and the varying degrees of reliable success across different emotion-types; that generative emotional competence must be accordingly localised. The question facing us at the end of this section, then, was *how* localised generative emotional competence ought to be, i.e. whether EC ought to advocate on behalf of a Type-Local or Scenario-Local interpretation of the scope of generative competence.

Accordingly, §6.5 endeavoured to provide a motivated answer to the question. Here, I argued that the defender of EC should refine their view in terms of Scenario-Local, i.e. that generative emotional competences ought to be assessed and localised with respect to emotion-types experienced in response to particular eliciting scenarios. In support of this, I suggested that the most plausible explanation of the way in which generative emotional competence develops is via mechanisms and processes of *emotional learning*. That is, in a way loosely analogous to the case of perceptual learning, subjects can come to be reliably successful in experiencing an emotion appropriately responsive to the presence of an evaluative property as a result of prolonged exposure to and repeated experience with a particular eliciting scenario. In the interest of advocating on behalf of a crystallised version of EC supplemented by this notion of emotional learning, I then proposed the Learned Emotional Competence view (LEC) of emotional justification, according to which S's emotional experience which attributes evaluative property E to object O is capable of immediately and defeasibly justifying her belief that O is E if and only if (i) S has the emotion as a result of her having emotionally learned that O is E, and (ii) S forms the belief that O is E on the basis of that emotion. §6.5.1 then sought to clarify the processes and epistemic significance of emotional learning, while §6.5.2 identified what I take to be a powerful selling point for LEC, namely, that it is uniquely capable of both accommodating for and providing a novel perspective on a popular insight from feminist philosophy, namely, that certain emotions experienced by members of socially and politically oppressed groups bear epistemic value. Finally, §6.6 strived to further strengthen LEC by explaining how the view meets the various desiderata identified earlier for a plausible theory of emotional justification. In what follows, I will conclude this dissertation by briefly summarising the key arguments made throughout the thesis and providing some final remarks.

## Conclusion

The purpose of this thesis has been to provide a thorough investigation into the question of *how* and *when* emotional experience is capable of lending immediate and defeasible epistemic justification to evaluative belief. I have argued that emotional experience is capable of performing such an epistemic role when and only when that emotional experience is formed via the subject's learned emotional competences to emotionally identify an object or state of affairs as instantiating the relevant evaluative property. In this last section, I will provide a brief review of the arguments made, and conclude with some final remarks.

### 7.1. Review

In Chapter 1, I began by introducing the justificatory thesis of emotion, and identifying support for this view within the contemporary philosophical literature. I then set out the parameters of the forthcoming discussion by commenting on some preliminary research questions concerning topics such as the nature of emotional experience, epistemic justification, and evaluative properties and belief.

In Chapter 2, I then presented *emotional dogmatism* as our first contender for a view of emotional justification. According to basic emotional dogmatism (ED), S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if that emotional experience makes it seem to her that *e*. It was argued that ED cannot be a plausible account of emotional justification given that it falls foul of the over-generalisation problem, i.e. ED ends up bestowing justificatory power to emotional seemings which *cannot* be capable of justifying the relevant evaluative belief. In light of this failure, then, I suggested that a *restricted* emotional dogmatist view (RED) based on Chudnoff's presentationalism is better placed to provide a plausible account of how emotional experience can justify belief. On this view, S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if that emotional experience (i) makes it seem to her that *e* and (ii) makes it seem as if she's emotionally aware of a truth-maker for *e*. I then attempted to bolster this view by defending it against a conspicuous objection facing it, namely, the objection concerning the nature of evaluative properties and the limits of emotional awareness. To this end, I appealed to a variety of sources and existing views in the literature in order to show that the defender of RED has the resources available to her to provide prima



facie plausible accounts as to how evaluative properties can be either direct or indirect objects of emotional seeming awareness.

In Chapter 3, I identified what I take to be the *real* problem facing RED. Namely, that it faces a troubling dilemma in what truth-makers for evaluative properties consist in. Either the truth-makers are the evaluative properties themselves, in which case RED continues to over-generalise, or the truth-makers are the non-evaluative properties which *give rise* to the evaluative properties that make the proposition true, in which case RED ends up *under-generalising* insofar very few emotional experiences bring the wide-ranging seeming awareness of non-evaluative properties required for them to possess justificatory power. In the wake of RED's failure, I then identified two alternative instantiations of a restricted emotional dogmatist view: (i) the McGrath-inspired Receptive Seemings Emotional Dogmatism (RSED), and (ii) the Markie-inspired Knowledge-How Emotional Dogmatism (KHED). I argued that both of these views fail in virtue of being unable to accommodate for the immediacy of the justification conferred by emotional experience. The overarching lesson of Chapters 2 and 3, I take it, is that emotional phenomenology and awareness are not suitable philosophical tools through which emotional experience can deliver immediate and defeasible epistemic justification to evaluative belief.

In Chapter 4, I then presented *emotional reliabilism* (ER) as the second contender for a view of emotional justification, according to which S's emotional experience is capable of immediately and defeasibly justifying her evaluative belief *e* if and only if S's forming evaluative beliefs on the basis of their emotional experiences is a reliable process of belief-formation. After motivating and defending ER against three preliminary objections, I then endeavoured to provide a defence of the view against the unreliability problem, i.e. the objection that forming evaluative beliefs on the basis of emotional experience is *not* a sufficiently reliable process of belief-formation. Here, I argued that all three of what I take to be the most compelling ways of arguing in favour of the emotions' systematic unreliability fail, such that, in lieu of a plausible alternative argument, we have no convincing theoretical support for the conclusion that forming evaluative beliefs on the basis of emotional experience is an unreliable process of belief-formation. ER, therefore, cannot be undermined by the unreliability problem *as is*.

In Chapter 5, I identified and presented what I take to be the *real* problem facing ER, namely, that emotional reliability cannot be sufficient for the presence of epistemic justification given the intuitive verdict delivered by an emotional analogue to traditional clairvoyance cases. I then considered two possible modifications of ER: (i) Evidential Emotional Reliabilism (EER) and (ii) Agent Emotional Reliabilism (AER). I rejected each of these in §5.3.1 and §5.3.2 respectively. I thereby concluded that ER cannot offer a plausible account as to how emotional experience can immediately and defeasibly justify evaluative belief. The overarching lesson of Chapters 4 and 5, I take it, is that emotional reliability, while perhaps important for epistemic justification, is not enough *by itself* to provide us with a satisfactory version of the justificatory thesis.

In Chapter 6, I presented the final category of views to be considered in this dissertation, namely, agent-based views of emotional justification. These views adopted a virtue epistemological framework in order to support the justificatory thesis of emotion, according to which features of the emoter *herself* are centrally relevant to determining whether her emotional experiences are capable of bearing justificatory power. Here, in light of virtue responsibilism's incompatibility with the justificatory thesis of emotion, and in light of virtue reliabilism's comparative compatibility, I identified Sosa's particular brand of virtue reliabilism as the most promising epistemological foundation for an agent-based view. I then presented Carter's Emotional Competence view (EC) as a version of the justificatory thesis of emotion grounded in Sosa's competence-based virtue reliabilism, and highlighted the central selling points of the account. I then argued, however, that EC must be specified further in order to bolster the view into the most compelling and plausible account it can be. To this end, I argued that EC is best supplemented with a learning-based framework for narrowly individuating emotional competences, such that this refined Learned Emotional Competence view (LEC) can capture all of the identified desiderata for a plausible view of emotional justification while explaining the distinct epistemic value of emotions experienced by members of oppressed social groups.

## 7.2. Summing Up

Emotional experience is a significant constituent of our mental lives. This is the claim that we began with in the dissertation, and that which initially motivated our investigation into the emotions' epistemic significance. Now, having reflected upon the arguments made here, and

having defended a virtue epistemological account as to how emotional experience can immediately and defeasibly justify evaluative belief, I hope that this significance is clearer than it was before we began. The emotions are not *just* epistemically significant because they can immediately and defeasibly justify evaluative belief, rather, more accurately, they are significant because of the way in which they are capable of doing it. They are significant because they are forged and moulded by long-term processes of learning and association; because when our emotional experiences *can* justify our evaluative beliefs, it is because, through time, we have developed learned competences to be reliably successful in experiencing emotions that are appropriately responsive to the presence of evaluative properties, and to be reliably successful at taking those experiences at face value.

The epistemic significance of emotions, then, is tied to the features of the emoter herself; to her history of past experiences and associations, and the ways in which her emotional dispositions have been shaped by that history. Views of emotional justification that neglect this broad holistic picture, as we've seen from the shortcomings of dogmatism and reliabilism, struggle to provide a nuanced and compelling theory as to how emotional experience can immediately and defeasibly justify evaluative belief.

## Bibliography

Alston, W. (2000) 'Virtue and Knowledge' *Philosophy and Phenomenological Research* 60(1) pp. 185 - 189

Annas, J. (2005) 'Comments on John Doris's Lack of Character' *Philosophy and Phenomenological Research* 71(3) pp. 636 – 642

Aristotle, *Nicomachean Ethics* in *The Complete Works of Aristotle*, Jonathan Barnes (ed.) (Princeton: Princeton University Press) (1984)

Audi, R. (1998) 'Moderate Intuitionism and the Epistemology of Moral Judgment' *Ethical Theory and Moral Practice* 1(1) pp. 15 – 44

--- (2013) *Moral Perception* (Princeton: Princeton University Press)

Badhwar, N.K. (1996) 'The Limited Unity of Virtue' *Noûs* 30(3) pp. 306 – 329

Baysan, U. (2017) 'A New Response to the New Evil Demon Problem' *Logos and Episteme* 8(1) pp. 41 - 45

Beebe, J. R. (2004) 'The Generality Problem, Statistical Relevance, and the Tri-Level Hypothesis' *Noûs* 38(1) pp. 177 - 195

Bergmann, M. (2006) *Justification Without Awareness: A Defence of Epistemic Externalism* (Oxford: Oxford University Press)

Bernecker, S. (2008) 'Agent Reliabilism and the Problem of Clairvoyance' *Philosophy and Phenomenological Research* 76(1) 164 – 172

Bonjour, L. (1980) 'Externalist Theories of Empirical Knowledge' *Midwest Studies in Philosophy* 5 pp. 53 – 73

- Borges, M. (2004) 'What Can Kant Teach Us about Emotions?' *The Journal of Philosophy* 101(3) pp. 140 – 158
- Brady, M. (2010) 'Virtue, Emotion, and Attention' *Metaphilosophy* 41(1 – 2) pp. 115 - 131
- (2013) *Emotional Insight: The Epistemic Role of Emotional Experience* (Oxford: Oxford University Press)
- Brentano, F. (1889) *The Origin of Our Knowledge of Right and Wrong*, Translated by Chisholm & Schneewind (London: Routledge and Kegan Paul, 1969)
- Brogaard, B. & Chudnoff, E. (2016) 'Against Emotional Dogmatism' *Philosophical Issues* 26(1) pp. 59 – 77
- Brogaard, B. & Gatzia, D. E. (2018) 'The Real Epistemic Significance of Perceptual Learning' *Inquiry* 61(5 – 6) pp. 543 - 558
- Carter, J. A. (2020) 'Epistemic Perceptualism, Skill, and the Regress Problem' *Philosophical Studies* 177 pp. 1229–1254
- Chudnoff, E. (2012) 'Presentational Phenomenology' in *Consciousness and Subjectivity*, Miguens & Preyer (eds.) pp. 51 – 72
- (2013) *Intuition* (Oxford: Oxford University Press)
- (2018) 'Epistemic Elitism and Other Minds' *Philosophy and Phenomenological Research* 96(2) pp. 276 – 298
- (2020) 'Experience and Epistemic Structure: Can Cognitive Penetration Result in Epistemic Downgrade?' in *Inference and Consciousness*, Chan & Nes (eds.)
- Comesaña, J. (2002) 'The Diagonal and the Demon' *Philosophical Studies* 110(3) pp. 249 – 266

- (2006) 'A Well-Founded Solution to the Generality Problem' *Philosophical Studies* 129 pp. 27 – 47
- (2010) 'Evidentialist Reliabilism' *Noûs* 44(4) pp. 571 – 600
- Conee, E. & Feldman, R. (1985) 'Evidentialism' *Philosophical Studies* 48(1) pp. 15 - 34
- (1998) 'The Generality Problem for Reliabilism' *Philosophical Studies* 89(1) pp. 1 – 29
- Cook, M. & Mineka, S. (1990) 'Selective Associations in the Observational Conditioning of Fear in Rhesus Monkeys' *Journal of Experimental Psychology: Animal Behaviour Processes* 16(4) pp. 372 - 389
- Cowan, R. (2016) 'Epistemic Perceptualism and Neo-Sentimentalist Objections' *Canadian Journal of Philosophy* 46(1) pp. 59 – 81
- (2018) 'Epistemic Sentimentalism and Epistemic Reason-Responsiveness' in *Evaluative Perception*, Bergqvist & Cowan (eds.) pp. 219 - 236
- D'Arms, J. (2005) 'Two Arguments for Sentimentalism' *Philosophical Issues* 15(1) pp. 1 – 21
- D'Arms, J. & Jacobsen, D. (2000) 'Sentiment and Value' *Ethics* 110(4) pp. 722 – 748
- (2003) 'The Significance of Recalcitrant Emotion' *Royal Institute of Philosophy Supplement* 52 pp. 127 - 145
- Darley, J. & Batson, C. (1973) 'From Jerusalem to Jericho: A Study of Situational and Dispositional Variables in Helping Behaviour' *Journal of Personality and Social Psychology* 27 pp. 100 - 108
- Deonna, J. & Teroni, F. (2012a) *The Emotions: A Philosophical Introduction* (New York: Routledge)

--- (2012b) 'From Justified Emotions to Justified Evaluative Judgments' *Dialogue* 51(2) pp. 55 - 77

De Sousa, R. (1987) *The Rationality of Emotion* (Cambridge, MA: MIT Press)

Dokic, J. & Lemaire, S. (2013) 'Are Emotions Perceptions of Value?' *Canadian Journal of Philosophy* 43(2) pp. 227 – 247

Döring, S. (2003) 'Explaining Action by Emotion' *The Philosophical Quarterly* 53(211) pp. 214 - 230

--- (2007) 'Seeing What to Do: Affective Perception and Rational Motivation' *Dialectica* 61(3) pp. 363 – 394

Doris, J. (1998) 'Persons, Situations, and Virtue Ethics' *Noûs* 32 pp. 504 – 530

Echeverri, S. (2019) 'Emotional Justification' *Philosophy and Phenomenological Research* 98(3) pp. 541 - 566

Elgin, C. (2008) 'Emotion and Understanding' in *Epistemology and Emotions*, Brun, Dogluoglu & Kuenzle (eds.) (Ashgate Publishing Company) pp. 33 – 50

Fantl, J. & McGrath, M. (2002) 'Evidence, Pragmatics, and Justification' *Philosophical Review* 111(1) pp. 67 – 94

Fricke, M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford: Oxford University Press)

Friedman, J. (2019) 'Inquiry and Belief' *Noûs* 53(2) pp. 296 - 315

Gallie, W. B. (1955) 'Essentially Contested Concepts' in *Proceedings of the Aristotelian Society* 56 pp. 167 – 198

- Gert, J. (2012) *Normative Bedrock: Response-Dependence, Rationality, and Reasons* (Oxford: Oxford University Press)
- Goldberg, S. (2010) *Relying on Others: An Essay in Epistemology* (New York: Oxford University Press)
- Goldie, P. (2004) 'Emotion, Feeling, and Knowledge of the World' in *Thinking About Feeling: Contemporary Philosophers on Emotions*, Robert C. Solomon (ed.) (New York: Oxford University Press) pp. 91 – 106
- Goldman, A. (1988) 'Strong and Weak Justification' *Philosophical Perspectives* 2 pp. 51 - 69
- (2008) 'Immediate Justification and Process Reliabilism' in *Epistemology: New Essays*, Quentin Smith (ed.) (Oxford: Oxford University Press) pp. 63 – 82
- Graver, M. (2007) *Stoicism & Emotion* (Chicago: The University of Chicago Press)
- Greco, J. (1999) 'Agent Reliabilism' *Philosophical Perspectives* 13 pp. 273 – 296
- (2003) 'Review: Further Thoughts on Agent Reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi' *Philosophy and Phenomenological Research* 66(2) pp. 466 - 480
- Harman, G. (1999) 'Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error', *Proceedings of the Aristotelian Society* 99 pp. 315 – 331
- Harrison, E. (2019) 'Seemings, Truth-Makers, and Epistemic Justification' *Synthese* (forthcoming) <https://doi.org/10.1007/s11229-019-02428-x>
- Hawthorne, J. & Stanley, J. (2008) 'Knowledge and Action' *Journal of Philosophy* 105(10) pp. 571 – 590
- Hoerl, C. (2001) 'The Phenomenology of Episodic Recall' in *Consciousness and Self-Consciousness*, Hoerl & McCormack (eds.) (Oxford: Clarendon Press) pp. 315 - 335



Hookway, C. (1994) 'Cognitive Virtues and Epistemic Evaluations' *International Journal of Philosophical Studies* 2(2) pp. 211 – 227

--- (2002) 'Emotions and Epistemic Evaluations' in *The Cognitive Basis of Science*, Carruthers, Siegal, and Stich (eds.) (Cambridge: Cambridge University Press) pp. 251 - 262

Huemer, M. (1999) 'The Problem of Memory Knowledge' *Pacific Philosophical Quarterly* 80(4) pp. 346 - 357

--- (2001) *Skepticism and the Veil of Perception* (Michigan: Rowman & Littlefield Publishers, Inc.)

---(2007) 'Compassionate Phenomenal Conservatism' *Philosophy and Phenomenological Research* 74(1) pp. 30 - 55

--- (2013) 'Phenomenal Conservatism Über Alles' in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Tucker (ed.) pp. 328 - 351

Isen, A. & Levin, P. (1972) 'Effect of Feeling Good on Helping: Cookies and Kindness' *Journal of Personality and Social Psychology* 21(3) pp. 384 – 388

Jaggar, A. (1989) 'Love and Knowledge: Emotion in Feminist Epistemology' *Inquiry* 32(2) pp. 151 – 176

James, W. (1884) 'What is an Emotion?' *Mind* 9(34) pp. 188 – 205

Johnston, M. (2001) 'The Authority of Affect' *Philosophy and Phenomenological Research* 63(1) pp. 181 - 214

Kauppinen, A. (2013) 'Ethics and Empirical Psychology' in *Empirically Informed Ethics*, Markus Christen (ed.) (Cham: Springer International Publishing) pp. 279 – 305

Kelp, C. (2019) 'How to be a Reliabilist' *Philosophy and Phenomenological Research* 98(2) pp. 346 - 374

Kornblith, H. (2009) 'A Reliabilist Solution to the Problem of Promiscuous Bootstrapping' *Analysis* 69 pp. 263 – 267

Lehrer, K. (2000) *A Theory of Knowledge* (2<sup>nd</sup> Ed.) (Boulder: Westview)

Lehrer, K. & Cohen, S. (1983) 'Justification, Truth, and Coherence' *Synthese* 55 pp. 191 - 207

Levin, D.T. & Banaji, M.R. (2006) 'Distortions in the Perceived Lightness of Faces: The Role of Race Categories' *Journal of Experimental Psychology: General* 135(4) pp. 501 – 512

Littlejohn, C. (2012) *Justification and the Truth-Connection* (Cambridge UK: Cambridge University Press)

Lyons, J. (2009) *Perception and Basic Beliefs. Zombies, Modules, and the Problem of the External World* (Oxford: Oxford University Press)

--- (2011) 'Circularity, Reliability and the Cognitive Penetrability of Perception' *Philosophical Issues* 21(1) pp. 289 – 311

--- (2013) 'Should Reliabilists Be Worried About Demon Worlds?' *Philosophy and Phenomenological Research* 86(1) pp. 1 - 40

--- (2019) 'Algorithm and Parameters: Solving the Generality Problem for Reliabilism' *Philosophical Review* 128(4) pp. 463 - 509

Macpherson, F. (2012) 'Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism' *Philosophy and Phenomenological Research* 84(1) pp. 24 – 62

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

- Markie, P. (2005) 'The Mystery of Direct Perceptual Justification' *Philosophical Studies* 126(3) pp. 347 – 373
- (2013) 'Searching for True Dogmatism' in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Tucker (ed.) pp. 248 – 269
- McDowell, J. (1985) 'Values and Secondary Qualities' in *Morality and Objectivity*, Honderich (ed.) pp. 110 – 129
- McGrath, M. (2013) 'Phenomenal Conservatism and Cognitive Penetration: The 'Bad Basis' Counterexamples' in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Tucker (ed.) pp. 225 - 247
- Milgram, S. (1974) *Obedience to Authority* (New York: Harper & Row)
- Milona, M. & Naar, H. (2020) 'Sentimental Perceptualism and the Challenge from Cognitive Bases' *Philosophical Studies* 177(10) pp. 3071 - 3096
- Mitchell, J. (2017) 'The Epistemology of Emotional Experience' *Dialectica* 71(1) pp. 57 – 84
- (2019) 'The Intentionality and Intelligibility of Moods' *European Journal of Philosophy* 27(1) pp. 118 – 135
- Montague, M. (2009) 'The Logic, Intentionality, and Phenomenology of Emotion' *Philosophical Studies* 145(2) pp. 171 – 192
- Moon, A. (2012a) 'Knowing Without Evidence' *Mind* 121(482) pp. 309 – 331
- (2012b) 'Three Forms of Internalism and the New Evil Demon Problem' *Episteme* 9(4) pp. 345 – 360
- Narayan, U. (1988) 'Working Together Across Difference: Some Considerations on Emotions and Political Practice' *Hypatia* 3(2) pp. 31 – 47

- Nussbaum, M. (2001) *Upheavals of Thought: The Intelligence of Emotions* (U.S.A: Cambridge University Press)
- Oddie, G. (2005) *Value, Desire, and Reality* (Oxford: Oxford University Press)
- O'Shaughnessy, B. (2003) *Consciousness and the World* (Oxford: Oxford University Press)
- Pelser, A. (2011) *Emotion, Evaluative Perception, and Epistemic Goods*, PhD Thesis, Baylor University
- (2014) 'Emotion, Evaluative Perception and Epistemic Justification' in *Emotion and Value*, Sabine Roeser and Cain Todd (eds.) (Oxford: Oxford University Press) pp. 107 – 123
- Perky, C. W. (1910) 'An Experimental Study of Imagination' *The American Journal of Psychology* 21(3) pp. 422 – 452
- Poellner, P. (2016) 'Phenomenology and the Perceptual Model of Emotion' *Proceedings of the Aristotelian Society* 116(3) pp. 261 – 288
- Pollak, S. D., Messner, M. Kistler, D. J., Cohn, J. F. (2009) 'Development of Perceptual Expertise in Emotion Recognition' *Cognition* 110(2) pp. 242 – 247
- Prinz, J. (2004) *Gut Reactions: A Perceptual Theory of the Emotions* (U.S.A: Oxford University Press)
- Pritchard, D. (2006) 'Review: *Evidentialism: Essays in Epistemology*' *The Philosophical Quarterly* 56(222) pp. 147 - 149
- (2009) 'Apt Performance and Epistemic Value' *Philosophical Studies* 143(3) pp. 407 - 416
- Pryor, J. (2004) 'What's Wrong with Moore's Argument?' *Philosophical Issues* 14(1) pp. 349 – 378

Roberts, R. C. (1988) 'What an Emotion is: A Sketch' *Philosophical Review* 97(2) pp. 183 – 209

Ross, W. D. (1930) *The Right and the Good* (U.S.A: Oxford University Press)

Sabini, J. & Silver, M. (2005) 'Lack of Character? Situationism Critiqued' *Ethics* 115(3) pp. 535 – 562

Shafer-Landau, R. (2003) *Moral Realism: A Defence* (Oxford: Oxford University Press)

Siegel, S. (2012) 'Cognitive Penetrability and Perceptual Justification' *Noûs* 46(2) pp. 201 – 222

Silins, N. (2005) 'Deception and Evidence' *Philosophical Perspectives* 19 pp. 375 - 404

Silva, L. L. (2020) 'The Epistemic Role of Outlaw Emotions' *Ergo* (forthcoming)

Simion, M. (2018) 'No Epistemic Norm for Action' *American Philosophical Quarterly* 55(3) pp. 231 – 238

Smithies, D. (2019) 'On the Global Ambitions of Phenomenal Conservatism' *Analytic Philosophy* 60(3) pp. 206-244

Solomon, R. C. (1988) 'On Emotions as Judgments' *American Philosophical Quarterly* 25(2) pp. 183 – 191

Sosa, E. (1992) 'Generic Reliabilism and Virtue Epistemology' *Philosophical Issues* 2 pp. 79 - 92

--- (2009) *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I* (Oxford: Oxford University Press)

--- (2010) 'How Competence Matters in Epistemology' *Philosophical Perspectives* 24(1) pp. 465 - 475

Sreenivasen, G. (2002) 'Errors About Errors: Virtue Theory and Trait Attribution' *Mind* 111(441) pp. 47 – 68

Stevenson, R., Oaten, M., Case, T., Repacholi, B. (2014) 'Is Disgust Prepared? A Preliminary Examination in Young Children' *The Journal of General Psychology* 141(4) pp. 326 – 347

Szigeti, A. (2015) 'Sentimentalism and Moral Dilemmas' *Dialectica* 69(1) pp. 1 - 22

Tappolet, C. (2016) *Emotion, Values, and Agency* (Oxford: Oxford University Press)

Teng, L. (2018) 'Is Phenomenal Force Sufficient for Immediate Perceptual Justification?' *Synthese* 195(2) pp. 637 – 656

Todd, C. (2014) 'Relatively Fitting Emotions and Apparently Objective Values' in *Emotion and Value*, Roeser and Todd (eds.) (Oxford: Oxford University Press) pp. 90 - 104

Tollhurst, W. E. (1998) 'Seemings' *American Philosophical Quarterly* 35(3) pp. 293 - 302

Tucker, C. (2010) 'Why Open-Minded People Should Endorse Dogmatism' *Philosophical Perspectives* 24(1) pp. 529 – 545

Vance, J. (2014) 'Emotion and the New Epistemic Challenge from Cognitive Penetrability' *Philosophical Studies* 162(2) pp. 257 – 283

Wilkowski, B. M. and Robinson, M. D. (2012) 'When Aggressive Individuals See the World More Accurately: The Case of Perceptual Sensitivity to Subtle Facial Expressions of Anger' *Personality and Social Psychology Bulletin* 38(4) pp. 540 – 553

Williamson, T. (1997) 'Knowledge as Evidence' *Mind* 106(424) pp. 717 - 742

Zagzebski, L. (1996) *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press)

--- (2000) 'Responses' *Philosophy and Phenomenological Research* 60(1) pp. 207 – 219

--- (2003) 'The Search for the Source of Epistemic Good' *Metaphilosophy* 34(1/2) pp. 12 - 28