

# **Essays on The Emotions**

**by Christina H. Dietz**



St. Cross College, University of Oxford

Submitted in September 2018

for a Doctorate of Philosophy in Philosophy

## Table of Contents

|                                                                    |           |
|--------------------------------------------------------------------|-----------|
| <b>ABSTRACT.....</b>                                               | <b>5</b>  |
| <b>ACKNOWLEDGEMENTS .....</b>                                      | <b>6</b>  |
| <b>ESSAYS ON THE EMOTIONS.....</b>                                 | <b>7</b>  |
| <b>CHAPTER ONE .....</b>                                           | <b>8</b>  |
| <b>INTRODUCTION .....</b>                                          | <b>9</b>  |
| 1.1 PRELIMINARY DISTINCTIONS.....                                  | 10        |
| <i>Propositional Attitudes .....</i>                               | <i>11</i> |
| <i>Noun Phrase Constructions .....</i>                             | <i>12</i> |
| <i>Explanatory Reasons and Personal Reasons.....</i>               | <i>14</i> |
| 1.2 STRUCTURAL GROUND CLEARING .....                               | 14        |
| <i>Discussion Guide .....</i>                                      | <i>24</i> |
| <b>CHAPTER TWO .....</b>                                           | <b>26</b> |
| <b>REASONS AND FACTIVE EMOTIONS.....</b>                           | <b>27</b> |
| 2.1 THE KNOWLEDGE REQUIREMENT .....                                | 28        |
| 2.2 FACTIVE EMOTIONAL STATE ASCRIPTIONS AND PERSONAL REASONS ..... | 30        |
| 2.3 EMOTIONAL FACTIVES AND PERCEPTUAL FACTIVES .....               | 33        |
| 2.4 THE PERCEPTUAL MODEL OF EMOTIONS.....                          | 36        |
| 2.5 CONCLUDING REMARKS.....                                        | 40        |
| <b>CHAPTER THREE.....</b>                                          | <b>42</b> |
| <b>EMOTIONS, EVIDENCE, AND SAFETY .....</b>                        | <b>43</b> |
| 3.1 EMOTIONS AS EVIDENCE .....                                     | 43        |
| 3.2 MICHAEL BRADY’S OBJECTIONS.....                                | 48        |

|                                                                       |           |
|-----------------------------------------------------------------------|-----------|
| <i>The Conclusive Reasons Objection</i> .....                         | 48        |
| <i>The Motivating Response Objection</i> .....                        | 49        |
| <i>The Proxy Reasons Objection</i> .....                              | 53        |
| <i>The Double Counting Argument</i> .....                             | 55        |
| 3.3 THE EMOTIONS AND SAFETY .....                                     | 58        |
| <i>Safety and Knowledge</i> .....                                     | 59        |
| <i>Emotions and Safe Belief-Forming Processes</i> .....               | 63        |
| <i>Empirical Support</i> .....                                        | 65        |
| 3.4 CONCLUDING REMARKS .....                                          | 72        |
| <b>CHAPTER FOUR</b> .....                                             | <b>73</b> |
| <b>CONDITIONAL EMOTIONS</b> .....                                     | <b>74</b> |
| 4.1 THE CHALLENGE .....                                               | 74        |
| 4.2 THE COGNITIVE EXPLANATION.....                                    | 78        |
| 4.3 SOME NATURAL LANGUAGE DATA.....                                   | 85        |
| <i>The First-Person/Third-Person Data</i> .....                       | 85        |
| <i>The 'Right-Now' Data</i> .....                                     | 86        |
| <i>Finite Conjunction</i> .....                                       | 87        |
| 4.4 CONCLUDING REMARKS .....                                          | 88        |
| Appendix.....                                                         | 89        |
| <b>CHAPTER FIVE</b> .....                                             | <b>93</b> |
| <b>DOXASTIC COGNITIVISM</b> .....                                     | <b>94</b> |
| 5.1 PRELIMINARY DISTINCTIONS .....                                    | 95        |
| 5.2 ANTI-INTELLECTUALISM.....                                         | 98        |
| 5.3 SOME VARIETIES OF DOXASTIC COGNITIVISM .....                      | 101       |
| 5.4 INCONCLUSIVE ARGUMENTS AGAINST MODERATE DOXASTIC COGNITIVISM..... | 106       |

|                                                    |            |
|----------------------------------------------------|------------|
| <i>Superficial Linguistic Argument 1</i> .....     | 106        |
| <i>Superficial Linguistic Argument 2</i> .....     | 108        |
| <i>The Problem of Recalcitrant Emotions</i> .....  | 109        |
| <i>The Specification Challenge</i> .....           | 111        |
| 5.5 A POSITIVE CASE FOR DOXASTIC COGNITIVISM ..... | 115        |
| 5.6 CONCLUDING REMARKS .....                       | 122        |
| <b>BIBLIOGRAPHY</b> .....                          | <b>124</b> |

### **Abstract**

In these essays, I hope to shed light on a range of basic philosophical issues concerning the emotions. To this end, I engage with many of the central ideas and questions found in philosophy of emotion over recent decades; Including, notably, the perceptual model of emotion, cognitivist theories of emotion, and questions concerning their intentional nature and epistemic role. This thesis is comprised of an introductory chapter which details the theoretical architecture that unifies this body of work, followed by four closely related essays within the philosophy of the emotions—each relating in some way to their intentional structure and epistemic significance. Among the salient influences are Timothy Williamson’s recent work in epistemology and Paul Grice’s work on the nature of reasons.

### **Acknowledgements**

I would like to thank John Hawthorne, Timothy Williamson, Ofra Magidor, and everyone else who supported me throughout the process of writing these essays. Thank you. I couldn't have done it without you.

# Essays on The Emotions

# Chapter One

## Introduction

This thesis is comprised of four closely related essays within the philosophy of the emotions—each relating in some way to their intentional structure and epistemic significance. There are three structural ideas about the emotions which have featured prominently in the literature in recent decades. One idea is the so-called ‘*cognitivist idea*’ which is paradigmatically presented as the idea that the emotions are or somehow integrally involve evaluative judgments (Solomon, 1988; Nussbaum, 2001; 2004). The second idea is the *perceptual model* which explores the ways that the emotions are analogous to perceptual states (Brady, 2008; 2013). The third idea is that emotions somehow always integrally involve an intentional object (Goldie, 200;2004). In this thesis, I extensively engage with each of these ideas. In Chapter 5, I reject judgment cognitivism and develop a more defensible version of cognitivism which I call ‘Doxastic Cognitivism’. In Chapter 2, the perceptual model is discussed and discarded on the basis of issues to do with the reason-sensitivity of the emotions. The objectual thesis is shown to have even less in its favor and its shortcomings are presented later in this chapter.

As well as developing alternatives to structural themes prominent in the literature, this work addresses other structural questions that have been somewhat neglected. Chapter 4 emphasizes the importance of the distinction between offline-emotions and online-emotions and puts it to use in solving an interesting philosophical puzzle. The category of offline-emotions has not yet received substantial attention but it is important to the development of a well-rounded model of the functional role of the emotions in our cognitive lives.

The relation between emotions and reasons has been touched on in the literature (Gordon, 1987) but to my mind, has been underexplored owing to a tendency to theorize about this

relation without first establishing a proper taxonomy of reasons. One of the aims of this thesis is to gain a more nuanced understanding of the relationship between reasons and the emotions (Chapter 2 sets out to accomplish this structural goal).

On the epistemic side, the question that is most actively pursued in the literature is whether the emotions can count as bona fide reasons for evaluative belief (Brady, 2013), and thus provide an epistemic foundation for ethics. I engage extensively with this epistemic question in chapter 3, arguing that the emotions can indeed play an evidential role of this sort. I also pursue some epistemic questions that have been less discussed in the recent literature. First, I ask whether emotions can be knowledge conducive without playing an evidential role. In this connection, I find the safety model of knowledge suggestive. I also turn to the empirical literature on the emotions, and in combination with a philosophical discussion of the safety model, I provide a detailed account of the various ways emotions can be knowledge-conducive without being evidential. In developing Doxastic Cognitivism (the thesis that emotions integrally involve belief), it also becomes natural to ask whether emotions themselves can constitute (as opposed to merely facilitate) pieces of knowledge (I address this question in chapter 5). Taken together, the material that follows constitutes both a critical and constructive project. It critically reflects many of the leading ideas in the literature about the epistemic and intentional nature of the emotions, and constructively offers a fresh and unified account.

## **1.1 Preliminary Distinctions**

One persistent theme throughout this work is that it is important to pay careful attention to the details of, and differences between, the various ways in which we talk about the emotions (in natural language) as a starting point for theorizing about them. In this connection, I would

like to begin by priming the reader with a number of structural observations that are important to the discussions to follow.

### **Propositional Attitudes**

Sometimes when we describe emotional states, our descriptions take an emotional verb (or adjective phrase) coupled with a that-clause complement (e.g. John is happy that he is in Paris). It is standard to think that constructions of this form claim that a certain relation – a propositional attitude—holds between an agent and a proposition (where the that-clause expresses the relevant proposition, and the emotional verb expresses the relevant relation). For now, sentences of this sort may be schematized as follows: S V's that P. In these cases, I will talk about emotions as having propositional content.

An attitude verb (V) is factive in case the inference from 'S V's that P' to 'P' is valid. Some standard examples of factive attitudes include 'knows-that', 'discovers-that', and 'understands-that'. Interestingly, some emotional verbs are factive. Thus, Jones cannot regret that he left the party unless it is true that he left the party (other examples of factive emotional verbs include 'happy-that', 'angry-that', 'surprised-that', 'astonished-that').

The factivity of certain emotional verbs is widely acknowledged by theorists such as Unger (1975), Gordon (1987), and Williamson (2000). However, Goldie complains that the factivity idea is, 'based in part on the idea that emotions must have propositions as objects' (Goldie; 2000, p. 22). Recall, however, that the claim that the verb 'anger' is factive is merely the claim that the inference from 'S is angry that P' to 'P' is valid. It bears emphasis that this claim does not require that any case of anger is a case where, for some P, one is angry that P. Indeed, that claim is not plausible and is actually in tension with the factivity thesis. For example, in a case where I am angry because I falsely believe that P, the factivity thesis requires

that it not be true that I am angry that P. (And while it is true that I believe that P, I am obviously not angry that I believe that P.) The factivity thesis discourages the view that in this case there is some proposition, P, such that I am angry that P. For now, let us say that a state is a non-factive episode of a factive state type (e.g. angry-that) if it is describable by a claim of the form, ‘S V’s because S thought that P’ but not truly describable by a claim of the form ‘S V’s that P’.<sup>1</sup>

Of course, many emotional verbs are not factive in the first place and, thus, do not entail the truth of their propositional complements. For example, ‘Lonnie is worried that he will stand out at High Table’ does not entail that Lonnie will stand out at High Table. (For one thing, the College Porter might turn Lonnie away at the door which, in turn, causes Lonnie to buy dinner from a kebab van instead.)

An interesting feature of factive emotional state ascriptions is that they also entail knowledge. That is, for factive emotional verbs, the inference from ‘S V’s that P’ to ‘S knows P’ is valid. Of course, the mere fact that an attitude is factive does not tell us that this inference is valid. Nevertheless, as Unger (1975), Gordon (1987), and Williamson (2000) have observed, there is a strong case to be made for positing a knowledge requirement on such verbs. (In chapter two, I explore the knowledge requirement in detail.)

### **Noun Phrase Constructions**

Sometimes, when we describe our emotional states, our descriptions take an emotional verb and a noun phrase complement (NP), or an emotional adjective coupled with a preposition and a noun phrase complement. Examples of the first sort can be seen in phrases like, ‘Mark envies Tim’ or ‘Jones hates John’. Examples of the second sort can be seen in phrases like ‘Jones feels guilty about the way he acted’ or ‘John is very proud of his daughter’. When a claim of the form,

---

<sup>1</sup> I will have more to say about this idea in section 1.2 and again, in chapter two.

‘S V’s NP’ is true, and the noun phrase refers to an object, it is standard to think that the emotional state has an emotional object (a material thing, an event, or some state of affairs) that the emotion is somehow about or ‘directed at’.<sup>2</sup>

It is worth noting that there are subtle differences between various prepositions as they combine with emotional adjectives. For example, ‘angry at’ seems to have a unique flavour. One might be angry *about* a multiple-choice test but this would not make it natural to say one is angry *at* the multiple-choice test. Thus, it seems that ‘angry-at’ and ‘angry-about’ express slightly different relations to emotional objects. I will not investigate these fine-grained differences here but it is good to bear in mind that they exist.

For certain emotional adjectives, the verb form is used to describe a relation whereby an agent causes the relevant emotion in others. For example, the verb ‘to anger’ describes the relation whereby the agent causes the emotion in others as in, ‘I angered him’.

It is also worth mentioning some emotional states are described by combining the verb ‘feel’ with an adjective or a noun. So, for example, we might ‘feel anger’, ‘feel angry’, ‘feel pity’, ‘feel hatred’. We should be careful not to make assumptions at the outset about the relation between, say, ‘is angry’ and ‘feels angry’. In general, ‘feels F’ does not entail ‘being F’ and ‘being F’ does not entail ‘feels F’. For example, consider the contrast between ‘being old’ and ‘feeling old’. One might feel old without being old, or be old without feeling old. If, for emotional states, there is an especially intimate relation between feeling V and being V, then it will call for careful argument. (For my purposes, we need not explore this point further here.)

---

<sup>2</sup> Of course, it is a further question what it takes for a claim of the form ‘S regrets NP’ to be true (where the V is a factive emotional verb). I will address this question in due course.

## Explanatory Reasons and Personal Reasons

There is an important distinction to be made between the *reasons why* one feels a particular emotion and one's *reasons for* having a particular emotion. Let us call the former, 'explanatory reasons' and let us call the latter, 'personal reasons.' Explanatory reasons explain why an agent has particular emotion, and these reasons do not always line up with an agent's reasons for being in that particular emotional state.<sup>3</sup>

Notice that personal reasons are typically described in the possessive-reason construction (possessive-reasons are usually described by a claim of the form 'my/his/her reason for Ving is P').<sup>4</sup> For example, consider a case where one of John's reasons for feeling sad is that he is eating dinner alone. In this example, John takes the fact that he is eating dinner alone as a reason he has for feeling sad, and this fact also explains why John feels sad. It is important to notice that there may be other reasons that explain John's sadness which are not available to John as reasons for him to feel sadness. For example, suppose, unbeknownst to John, he has a neural imbalance which causes him to feel sadness. The fact that John has a neural imbalance which causes him to feel sad is a reason *why* John is sad but this fact is not among one of John's personal reasons for feeling sad. In sum, personal reasons tend to be explanatory but explanatory reasons are not always among an agent's personal reasons for Ving.<sup>5</sup> This distinction will be especially important in chapter two but will also be a structural theme in chapters three, four, and five.

### 1.2 Structural Ground Clearing

I have introduced two primary uses of emotional verbs. One use involves combining an emotional verb with a noun phrase complement, and the other use involves combining an

---

<sup>3</sup> Here, I follow Grice's terminology for reasons introduced in his, *Aspects of Reason* (2001).

<sup>4</sup> Jennifer Hornsby (2008) uses the 'reasons for which' construction as her way of picking out personal reasons. Hornsby glosses personal reasons as the *reasons for which* one acts.

<sup>5</sup> See Dietz (2016) for a more in-depth discussion of actions, reasons, and causes.

emotional verb with a that-clause complement. But how do these two constructions relate to each other? Before we set about theorizing about the emotions in a serious way, it is crucial to be clear about the two kinds of relations picked out by these constructions. As I see it, a good deal of the literature on emotions has suffered from not being more precise about this issue.

Emotion-describing verbs express different relations according to whether they take a that-clause or noun-phrase complement.<sup>6</sup> The most powerful case for the ambiguity of emotion-describing verbs turns on what happens when we substitute a noun phrase that picks out a proposition for a that-clause that expresses it. The result has entirely different truth conditions. Thus, consider the pair:

- i. I fear that the tiger will attack.
- ii. I fear the proposition that the tiger will attack.

In a context where the (i) is true, (ii) is almost certainly false – only a very strange individual would fear propositions. And it is very natural to explain the difference in truth conditions by supposing that ‘fear’ expresses different relations in each case. Let us begin by considering, ‘fear’. I will refer to the relation picked out by ‘fear’ in (i) as the P-fearing relation ( $\text{fearing}_P$ ) and the relation picked out by ‘fear’ in (ii) as O-fearing relation ( $\text{fearing}_O$ ). In this chapter, I will adopt the following conventions with respect to any emotional verb (V):

$V_O$  names the relation described by an emotional verb that takes a noun phrase (that picks out an emotional object) as its complement (e.g. John is afraid of Jones).

$V_P$  names the relation described by an emotional verb that takes a that-clause (which picks out a proposition) as its complement (e.g. John is afraid that Jones will attack him).

---

<sup>6</sup> In this way, they are analogous to ‘know’, which is well-known to be ambiguous between two different relations (roughly according to whether it takes a that-clause or noun phrase complement).

We can now ask the following question: *For any  $V$ , what is the relationship between the  $V_O$  relation and the  $V_P$  relation?* Let us begin with the following conjecture:

**Conjecture 1:** Whenever S stands in the  $V_P$  relation to a proposition, and x is the main subject of that proposition, S stands in the  $V_O$  relation to x.

Conjecture 1 is motivated by the following kind of reasoning: ‘If Marlowe is afraid that spider will bite, then surely, it is natural to say that Marlowe is afraid of the spider!’ Conjecture 1 produces terrible results. For example, Conjecture 1 predicts that if I am afraid that my house will burn down, then an x that I am afraid of is my house. Consider a weaker conjecture.

**Conjecture 2:** Whenever S stands in the  $V_P$  relation to some proposition, there is some object x that S stands in the  $V_O$  relation to.

We can motivate Conjecture 2 with the natural thought that when, for example, I fear that I will embarrass myself, the emotional object of my fear<sub>P</sub> is the event my *embarrassing myself*. The observation that Conjecture 1 fails (because I do not fear myself) is no problem for Conjecture 2. However, recall that for non-factive  $V_S$ , when S stands in the  $V_P$  relation to a proposition, this does not entail that P is true. ‘Fear’ is a paradigm example of a non-factive emotional verb because it is obvious that it can be true that *I fear that I will embarrass myself* even if it is false that I will. In the event where I do not embarrass myself, we might wonder what the emotional object of my fear<sub>O</sub> was. We cannot say it is the event of my embarrassing myself (since no such event occurred, no such event exists). In this case, it seems right to say that there is no emotional object that is feared<sub>O</sub>.

Thus, for non-factive emotional verbs like ‘fear’, the fact that S stands in the  $V_P$  relation to some proposition does not entail that S stands in the  $V_O$  relation to some object. For even if the fear has some propositional content, it is not true that there is always an associated emotional object  $x$ , such that  $S V_s x$ .

What about the case of factive verbs? Conjecture 1 does not fare any better when restricted to factive emotional verbs. For example, I might regret that my party ended so soon but it does not follow that I regret my party.

In the case of factive emotional verbs, Conjecture 2 seems compelling. The recipe for constructing counterexamples to Conjecture 2 does not work for factive emotional verbs.<sup>7</sup> This is because the truth of the ‘S V’s that P’ entails P for factive verbs. As such, if some event is required for truth of P, the truth of ‘S V’s that P’ guarantees that the event will occur (or has occurred). For example, the truth of ‘John regrets that he placed the bet’ guarantees that there was the event of *John’s placing the bet*, and such an event makes a perfectly appropriate object for a  $V_O$  claim: John regrets *the placing of the bet*. This phenomenon generalizes. If I regret<sub>P</sub> that I ordered scallops for dinner, I regret<sub>O</sub> the event of my *ordering scallops*. In the case of factive emotional verbs, one can also appeal to the fact as the emotional object. For example, suppose that Mark is surprised that no one applauded him after his talk. In this case, there is no material object for Mark to be surprised at but there is an event (of *his getting no applause*) and also a fact (about the event) that Mark is surprised at.<sup>8</sup> In such a case, it seems perfectly felicitous to say that Mark is *surprised at the fact* that no one applauded him after his talk.

---

<sup>7</sup> The recipe: Find a case where the event—whose occurrence is required for the truth of the proposition—does not occur.

<sup>8</sup> Goldie (2000) recognizes that in some cases, the object of an emotion will be an event (not a material thing) (see p. 17). One subtle point that I will not dwell on: In a case where one is angry about an event (say, the death of a loved one) it is a little more natural (at least in some idiolects) to say that one is angry *about* the death, rather than one is angry *at* the death. Even more obviously, if one is angry at John, it is

In sum, for factive emotional verbs, if S V's that P, then there will always be some associated emotional object that S can be said to V<sub>O</sub>. Thus, the following conjecture seems correct:

**Conjecture 3:** For any factive emotional V, whenever S stands in the V<sub>P</sub> relation to some proposition, there is some object x that S stands in the V<sub>O</sub> relation to.

It seems clear enough that Conjecture 3 cannot be extended to non-factive emotional verbs. Thus, slogans like, 'all emotions have emotional objects' or 'all emotions are directed at objects in the world' are certainly false when we consider the fact that non-factive emotional verbs do not guarantee emotional objects. Indeed, this consideration is why Conjecture 2 is false. Yet, it is worth mentioning that some philosophers argue that there are nonexistent objects (following Meinong). Although it is a minority view, this approach offers a unique way of saving Conjecture 2. The Meinongian can save Conjecture 2 by claiming that in cases of non-factive emotional verbs, there is an emotional object that S stands in the V<sub>O</sub> relation to—namely, a non-existent emotional object. In the case where I fear that I will embarrass myself at the party and I do not, the Meinongian will say that the emotional object that I fear<sub>O</sub> is the non-existent-occurrence of my *embarrassing myself*.<sup>9</sup>

Goldie (2000) argues for the slogan that all 'emotions are directed towards objects in the world.'<sup>10</sup> Goldie's 'feeling towards' account is quite Meinongian inasmuch as it allows for non-

---

far less natural to say that one is angry about John. 'Angry about' seems to work more naturally with noun phrases that do not pick out people.

<sup>9</sup> Meinongian ways of describing the world sound truly bizarre to me. A Meinongian would say things like, "Oh, no! There is such a thing as x, it just doesn't exist." Russell writes that Meinong's view is, 'intolerable; and if any theory can be found to avoid this result, it is surely to be preferred.' (1905, p. 483)

<sup>10</sup> See Goldie (2000). In chapter two, Goldie discusses what is involved in emotional experiences. He writes, 'our entire mind and body is engaged in the emotional experience, and all the feelings are 'united in consciousness' in being directed towards its object: united 'body and soul', 'heart and mind'. For

existent emotional objects even though orthodoxy holds that the class of ‘objects in the world’ does not include non-existent objects. Here is not the place to evaluate Meinongianism, though I am not drawn to it myself. I should point out that even if one held it, the contrast suggested above between factive and non-factive verbs does not disappear. We can reformulate it as follows: For factive but not for non-factive emotional verbs, if S stands in the  $V_P$  relation (associated with the verb) to some proposition, S stands in  $V_O$  relation (associated with the verb) to some *existing* object.

There is another related contrast between factive and non-factive emotional verbs—namely, that for non-factive emotional verbs, claims of the form ‘S V’s NP’ can be true regardless of whether there is an object in the world picked out by the noun phrase. For example, one may truly fear God’s wrath even if there is no such thing as ‘God’s wrath’ that exists. Let us call this phenomenon, ‘Empty-Ving’.

**Empty-Ving:** A claim of the form ‘S Vs NP’ is true even though ‘NP’ does not pick out an object.

Somewhat tentatively, I would like to suggest that while the phenomenon of Empty-Ving certainly arises for non-factive emotional verbs, it does not seem to arise for factive emotional verbs. For example, it is hard to see how the proposition expressed by ‘John regrets his third birthday party’ is true if, in fact, John never had a third birthday party. Of course, John might have a feeling of regret prompted by a false belief that he had a third birthday party but in this case, the ascription is clearly false. If this right, we have a support for a further conjecture:

---

example, sexual desire is felt with the whole being body and soul for the one we desire.’ See also Goldie (2002).

**Conjecture 4:** Non-factive emotional verbs are subject to Empty-Ving but factive emotional verbs are not.

If Conjecture 4 is right, we can make a related observation: Factive emotional verbs (like ‘surprise’) are subject to exportation in a way that non-factive emotional verbs (like ‘fear’) are not.

**The Exportation Schema:** For emotional verbs, if S Vs NP then there was/is/or will be an object x that is identical to the NP such that S Vs x.

Thus, if I fear the present King of France, it does not follow that there is something identical to the present King of France that I fear (and this is explained by the phenomenon of Empty-Ving). But if I regret the murder, it does follow that there is something identical to the murder that I regret (assuming that there are not cases of Empty-Ving for factive verbs like ‘regret’). In sum, Conjecture 4 supports a further conjecture:

**Conjecture 5:** Factive Vs conform to the exportation schema.

Notice that because of Empty-Ving, exportation is a better test for the presence of a genuine emotional object than the mere use of a noun phrase. If there exists a God that I fear, then that is

a sure sign that there is an emotional object of my fear. But the truth of, ‘I fear God’ is not a sure sign of an emotional object since this may be true whether or not God exists.<sup>11 12</sup>

Let us now turn to another question: For factive and non-factive emotional verbs, does standing in the  $V_O$  relation to a genuine emotional object guarantee that one also stands in the  $V_P$  relation to some proposition? I will provide two arguments for the following answer:

**Conjecture 6:** It is not the case that, *for any emotional verb*, standing in the  $V_O$  relation (associated with the verb) to a genuine emotional object guarantees that there is some proposition such that one stands in the  $V_P$  relation (associated with the verb) to some proposition.

First consider factive emotional verbs. If John regrets the party ( $\text{Regret}_O$ ), is there any guarantee that there is some proposition  $P$  such that John regrets that  $P$  (where the main subject of  $P$  is the party)? Here are two reasons for thinking not: First, an agent can forget their personal reasons for feeling a particular emotion. For example, suppose Marlowe is angry that Frank popped her balloon and consequently, Marlowe is angry at Frank. Marlowe might forget the fact that

---

<sup>11</sup> Of course, explaining the semantics of these sentences is a delicate matter. The issues are similar to those that arise for ‘She wants the perfect husband’ or ‘She’s searching for the perfect husband’ which may be true claims despite the fact that the noun phrases do not designate an existing object. A great deal has been written on this topic. I will not explore the relevant complexities any further here. For further discussion see Partee, (1974) and Zimmerman (1993). We will likely need to proliferate relations beyond the  $V_P$  and  $V_O$  relations.

<sup>12</sup> Notice again that Meinongianism will make a real difference to the discussion. What orthodoxy takes as a case of Empty-Ving, the Meinongian will take as a case of a noun-phrase picking out a non-existent object. And where orthodoxy sees a non-exportable case, the Meinongian will see the case to be exportable (so long as one is careful not to claim existence). Thus, the Meinongian will endorse the inference from ‘John fears the monster’ to ‘There is something that is the monster (albeit something that may not exist) that John fears’. Although I do not wish to pursue Meinongianism any further here, it bears emphasis that Meinongianism does not undermine the important contrasts presented above. For example, the Meinongian is free to endorse ‘Quasi-Empty-Ving’ as an alternative to Empty-Ving and the contrast between factive and non-factive verbs would remain intact: Quasi-Empty-Ving: A claim of the form ‘ $S$  Vs  $NP$ ’ is true when ‘ $NP$ ’ does not pick out an existing object.

prompted her anger but still feel angry at Frank.<sup>13</sup> It is perfectly easy to imagine a state of directed-anger without an awareness of why one is angry.

Second, Marlowe's anger might be caused by a false belief. Suppose instead that Marlowe falsely believes that Frank popped her balloon. Because of factivity, it is wrong to say that Marlowe is angry that Frank popped her balloon.<sup>14</sup> And it is also wrong to say that Marlowe is angry that she believes that Frank popped her balloon. Here, we have a case of objectual anger<sub>O</sub> without there being any proposition available for us to say (truly), 'Marlowe is angry that P'.

The second argument above will not generalize to non-factive emotions since the argument above relies on the factivity of 'is angry that'. But the first argument does seem to generalize. In summary, Conjecture 6 seems to be correct even when restricted to either factive or non-factive verbs.

As we move forward, the contrast between V<sub>O</sub> relation and V<sub>P</sub> relation (as well as the complexity of the relations between them) will be important to bear in mind. Philosophers are not always so careful here. As an example, consider Tim Crane on the intentionality of sensations and emotions.<sup>15</sup> Crane uses 'intentional content' as a catchall term. Moreover, he lacks the resources to make sense of the ambiguity between expressing V<sub>O</sub> and V<sub>P</sub> relations. Granted, Crane distinguishes states of, say, fear that have propositions as contents and states that have other contents. But it is a mistake to think that this comes to the same thing as the distinction between fearing<sub>O</sub> and fearing<sub>P</sub>. Consider the following pair:

---

<sup>13</sup> The connection between emotions and personal reasons explored in chapter two will reinforce this judgment. In chapter two, I will suggest that insofar as anyone is angry that P at some time, P has to be one of their personal reasons for being angry at that time.

<sup>14</sup> See chapter two for a full account of the connection between factive emotional verbs, knowledge, and reasons. One will find an in-depth discussion of motivational reasons and error-cases in Alvarez (2010).

<sup>15</sup> See Tim Crane (2001; 2003).

- (i) I fear that I will never love again.
- (ii) I fear the proposition that I will never love again.

In terms of Crane's distinction, these are both cases where the content of the mental state is something that can be either true or false because both of the contents are propositions. But in terms of my distinction, only (i) is an expression of the  $V_P$  relation because the attitude verb takes a complementizer phrase. The kind of relation to a proposition expressed by 'John fears that he will lose' is very different to the relation to the very same proposition expressed by 'John fears the proposition that he will lose.' The distinction between these two kinds of relations is crucial to understanding the intentional structure of emotions.

Here is another example of confusion between  $V_O$  and  $V_P$  (one briefly alluded to earlier). About factive emotions, Goldie (2000) writes, 'Again, this is a sort of over-intellectualization of the emotions. The idea is based in part on the idea that emotions must have propositions as their objects' (p. 22) But here, Goldie is confused. We should all agree that when John is amazed that his books caught on fire, John is unlikely to be amazed<sub>O</sub> at the proposition that his books are on fire. Of course, it is always possible that in some eccentric emotional episode, a person is amazed at a proposition. (For example, maybe Jeff King could be amazed at a certain proposition because it provides an unexpected counterexample to some theory of propositions that he likes.) But such an emotional episode would be a case of amazement<sub>O</sub> and the factivity thesis only concerns emotional verbs which take that-clauses as complements and thus, only concerns the  $V_P$  relation picked out by verbs in such contexts.

The structure articulated in the above discussion is not incidental to the four essays that follow. Indeed, the discussion above reflects the architecture that I have brought to each of the

topics under investigation in this body of work. It was, thus, very important that I made this architecture vivid at the outset.

It is worth noting that the essays to follow are self-standing. As such, each chapter can be read without reading any of the others. That said, I believe that the chapters complement each other and cumulatively provide a vision of the emotions that speaks to a wide range of central issues concerning their intentional structure and epistemic roles.

### **Discussion Guide**

In chapter two, ‘Reasons and Factive Emotions’, I present and explore some ideas about how factive emotional states and factive perceptual states each relate to knowledge and reasons.<sup>16</sup>

In chapter three, ‘Emotions, Evidence, and Safety’, I explore two ways that the emotions can be of epistemic value with respect to evaluative beliefs. First, emotions can provide bona fide reasons to believe evaluative propositions, and thus, provide evidence for those propositions by raising the epistemic probability of their being true. Second, when emotions are not playing an evidential role, they can enhance the safety of evaluative belief formation and thus, facilitate knowledge.

In chapter four, ‘Conditional Emotions’, I present and solve a puzzle presented by certain embeddings of emotional verbs in conditionals (e.g. If John is at the pub, I’m surprised). Against competing solutions to the puzzle, I present a solution that turns on the nature of reasoning under supposition.

---

<sup>16</sup> A version of this paper was published in *Philosophical Studies* in 2017.

In chapter five, ‘Doxastic Cognitivism’, I explore and defend the cognitivist view that emotions involve beliefs. My defensive strategy involves carefully distinguishing belief from judgment and motivating an anti-intellectualist conception of belief itself.

## Chapter Two

## Reasons and Factive Emotions

As a number of authors have observed, certain emotional state ascriptions are factive.<sup>1</sup> In particular, certain state ascriptions of the form ‘S V’s that P’ entail P (where S names an agent and V stands for an emotional verb or adjectival phrase of the form ‘is F’).<sup>2</sup> Some examples of factive emotional state ascriptions are as follows: Jason regrets that he danced at the party. Ann is angry that she lost the bet. Lonnie is delighted that his mother ironed his pants.<sup>3</sup> By contrast, non-factive emotional state ascriptions (e.g. fears-that, worries-that, hopes-that, etc.) have no such entailment.<sup>4</sup> Factive emotional state ascriptions form a subclass of factive mental state ascriptions that include ‘knows that’, ‘sees that’ and ‘understands that’. In this discussion, I will explore some important similarities and differences between factive emotional state ascriptions

---

<sup>1</sup> See Gordon (1987), Unger (1975), and Williamson (2000).

<sup>2</sup> In the case of ‘happy’ there are other that-clause constructions that can be formed by combining a so called ‘linking verb’ with ‘happy’ followed by a that-clause: felt happy that P, seemed happy that P, remained happy that P, and so on. As far as any emotional adjective (F) is concerned, I will limit my focus to ‘is F that’ (and ‘was F that’).

<sup>3</sup> There are other factive environments in which emotional verbs can figure which will not be of interest to me. For example, ‘S is angry because P’ plausibly entails P (on account of the fact that ‘because’ induces a factive environment). I limit my focus to factive ‘V-that’ constructions here. I shall also not be discussing constructions which take the form of an emotional verb followed by a noun phrase. (e.g. Ann regrets the picnic) which are not of the right form to be factive.

<sup>4</sup> An alternative hypothesis is that a claim of the form ‘S is angry that P’ cancellably implicates P without entailing P. This is unpromising. Consider the speech: ‘She was angry that her house had burnt down but it hadn’t, it’s just that the radio reported that it had’. That speech is very marginal but would be fine if ‘she was angry that her house burnt down’ merely generated a cancellable implicature that her house had burnt down in the same way that standard implicature cancellation cases work. But this case is very different from the standard implicature cancellation cases and the contrast is striking. ‘There is a gas station around the corner but I don’t know whether it’s open’ is fine. ‘She is angry that her money has been stolen but doesn’t know that her money has been stolen’ is markedly worse. In general, ‘angry that’ patterns in relevantly similar ways to ‘know that’ as far as factivity goes. ‘She knew that her house burnt down but it hadn’t, it’s just that the radio reported that it had’ is marginal in the same way as the ‘angry that’ example. An alternative hypothesis to the cancellable implicature hypothesis is that ‘S is angry that P’ non-cancellably presupposes P without entailing P. But given the parallels between ‘knows that’ and ‘angry that’, I see little promise in a view that concedes factivity for ‘know that’ but opts for some fancy alternative to factivity for ‘angry that’.

and factive perceptual state ascriptions. This exercise will serve to illuminate some important structural features of factive emotional states—namely, how these states relate to reasons and knowledge—and in turn, will shed light on the so-called ‘perceptual model’ of the emotions (a view which has received considerable attention in the recent literature on emotions).

In 2.1, I present a brief defense of a knowledge requirement on factive emotional states.<sup>5</sup> In 2.2, I explain the source of the knowledge requirement for those states. This account will crucially depend on the connection between those states and what Grice calls ‘personal reasons’. In 2.3, I investigate the structural parallels and disanalogies between factive emotional state ascriptions and factive perceptual state ascriptions. In 2.4, I relate the previous discussion to the perceptual model of emotions.

## 2.1 The Knowledge Requirement

One fundamental assumption of the sections that follow is that there is a knowledge requirement on factive emotional state ascriptions:

**The Knowledge Requirement:** For factive emotional state ascriptions, if S V’s that P then S knows that P.

This claim is not new and is most famously defended by Unger (1975), Gordon (1987), and Williamson (2000). Of course, there are alternative hypotheses that respect the factivity of the state-ascriptions we are interested in but are weaker than the knowledge requirement. For example:

---

<sup>5</sup> I will sometimes use the term ‘factive emotional states’ as shorthand for emotional states described by factive emotional state ascriptions. Note that I am not going to pursue the metaphysical question as to whether the mental state token associated with ‘she is angry that she lost’ is essentially factive. This question turns on murky metaphysical issues about events/states (such as whether it is essential to that very anger event that it occurs in a world where she lost).

**True Belief Condition:** For factive emotional state ascriptions, if S V's that P then S has a true belief that P.

**Justified True Belief Condition:** For factive emotional state ascriptions, if S V's that P then S has a justified true belief that P.

However, I take it that Unger et al. are right to insist on the much stronger knowledge requirement on these states. The most direct argument in favor of the knowledge requirement on factive emotional state ascriptions is that the result of conjoining factive emotion ascription with knowledge denial is linguistically bizarre. Consider how strange it sounds to say, 'Sarah regrets that she punched Jason but Sarah does not know that she punched Jason'.<sup>6</sup>

Further, a wide range of judgments can be explained by the knowledge requirement. Suppose Maria sees Sarah's identical twin fall off a cliff and feels surprise. Here, Maria justifiably mistakes Sarah's twin for Sarah. Meanwhile, at the same time as the twin's accident, Sarah falls off a cliff too. In a case like this, Maria's belief that Sarah has fallen off a cliff is true and justified yet it would be infelicitous to say that Maria is surprised that Sarah fell off a cliff. This infelicity can be explained by the fact that Maria did not know that Sarah fell off a cliff. Notice that this is a case of justified true belief without knowledge and so the infelicity cannot be explained by recourse to either the true belief requirement or the justified true belief requirement. Similar remarks apply mutatis mutandis in cases where Sarah's identical twin prompts Maria's delight, anger, sadness, and so on. In the relevant scenarios, it is infelicitous to assert that Maria

---

<sup>6</sup> In correspondence, Mark Schroeder objected. He pointed out that, 'I believe that P but I don't know that P' also sounds weird but is not good evidence that belief entails knowledge. Schroeder's objection is confused. His example relies on the use of the first-person pronoun. Described in the third-person, we find no oddity at all: 'He believes that P but he doesn't know that P' is perfectly felicitous.

is delighted that Sarah fell off a cliff, or sad that Sarah fell off a cliff. The knowledge requirement can elegantly explain these infelicities.<sup>7 8</sup>

(Note that the knowledge requirement, as it applies to the verb ‘surprise’, concerns constructions of the form ‘S is surprised that P’. In the case just described, it would be more felicitous to say, ‘Maria is surprised at the thought that Sarah fell off a cliff’ or ‘Maria is surprised because she believes that Sarah fell off a cliff’ but these are different constructions. The knowledge requirement as stated should not be conflated with the grotesque overgeneralization that surprise requires knowledge).

Thus, ends my brief defense of the knowledge requirement on factive emotional state ascriptions. Adding substantially new arguments for the knowledge requirement on factive emotional states is not my task here. In the next section, I will explore the source of the knowledge requirement on such states. This will serve as a nice starting point for uncovering the structural contrasts between factive emotional states and factive perceptual states.

## **2.2 Factive Emotional State Ascriptions and Personal Reasons**

In this section, I will explain why the knowledge requirement holds for factive emotional states. But first let me say some things about my use of the Gricean expression, ‘personal reasons’. In connection with the emotions, reasons talk can take a variety of forms. When we are

---

<sup>7</sup> A commonly recycled objection to Williamson’s view about stative factives is that in fake barn cases, knowledge ascriptions are infelicitous but ascriptions of various factives such as, ‘She sees that there is a barn in front of her’, ‘She is happy that there is a barn in front of her’ etc. sound *prima facie* more acceptable. [This objection appears in Turri (2010), Hughes (2014) and Locke (2015)]. See Hawthorne and Magidor (forthcoming) for an extensive discussion of why this objection is unconvincing.

<sup>8</sup> One reader suggested that this data could be explained by a causal requirement instead of the knowledge requirement. The causal requirement would roughly be something like, ‘V-ing that P requires a true belief that P that is caused by the fact that P’. This idea doesn’t fare well when we turn to cases where one knows P but where one’s true belief is not caused by the fact that P (the most obvious case would be one where P is a fact about the future.) The absence of a causal link does not block factive emotion ascriptions in such cases. For relevant discussion, see Dietz (2016).

interested in reasons that can explain an agent's anger, we typically make use of a 'reason why' construction. (The reason why S is angry is that P). When we are interested in the normative reasons available to an agent for feeling anger, we typically make use of a 'reason to be angry' construction. (P is a reason for S to be angry). But for my purposes here, I am primarily interested the kind of reasons that are picked out by paradigmatic uses of possessive constructions of the form 'my/his/her reason for V-ing is...' Following Grice (2001), I use the expression, 'personal reasons' to refer to those reasons standardly picked out by this kind of possessive construction.<sup>9</sup> (Many philosophers use the expression, 'motivating reasons' here). As Grice notes, there are two kinds of constructions used to describe personal reasons, that-clause constructions and infinitival constructions: (i) His reason for leaving was that the house was on fire. (ii) His reason for leaving was to avoid harm.

I am interested in those personal reasons described by expressions which contain that-clauses as exemplified by (i). My use of the expression 'personal reason' will hereafter refer to kind (i). With the ideology of personal reasons in place, consider two conjectures:

**Conjecture One:** For factive emotional state ascriptions, if S V's that P then one of S's personal reasons for V-ing is P.<sup>10</sup>

---

<sup>9</sup> Given that the possessive construction is notoriously context-sensitive, we should expect there to be other uses of the 'my/his/her reason' construction. As Hornsby (2008) notes, the kind of relationship expressed by 'her reason' in 'her reason for going to the psychiatrist is that she thinks she is being followed' and 'her reason for running is that she is being followed' is different to that in 'her reason for running is that she thinks she is being followed.' As Hornsby would put it, the sentence following 'that' in the final example does not express the 'reason for which' one acts. It is the 'reason for which' use of the possessive reason construction that interests me here.

<sup>10</sup> The reverse of conjecture one—if P is one of S's reasons for V-ing, then S V's that P—is far less plausible. One concern has to do with certain cases where P is a fact about the agent's psychological state. If one of John's reasons for being angry is that John thinks that his friend has lied, it is not plausible that John is angry that he thinks that his friend has lied. (But this may involve a different use of the possessive reason construction, see footnote 9.) But troublesome cases are not limited to cases that involve psychological propositions as the complement of the reason construction. If Margot's car has been stolen, one of Margot's reasons for being angry might be that her car is brand new. But in this case, it doesn't sound right to say that Margot is angry that her car is brand new. This raises the interesting

**Conjecture Two:** If one of S's personal reasons for V-ing is P, then S knows P.<sup>11</sup>

These two conjectures entail the knowledge requirement and if true, offer a satisfying explanation of that requirement.

Conjecture one is extremely plausible. It would be strange to say that John is angry that his treasure has been stolen but that the proposition that his treasure has been stolen is not one of John's reasons for being angry.

Conjecture two suggests that knowledge is a condition on having personal reasons and this claim is also very plausible. Of course, we should have no expectation that there is a knowledge condition on explanatory non-personal reasons. If an explanatory reason why Jason feels jealous is that he is chemically imbalanced, then the fact that he is chemically imbalanced isn't necessarily available to Jason as a reason for him to feel jealous. In a case like this, it would be out of place to say that Jason's reason for feeling jealous is that he is chemically imbalanced. Explanatory reasons that are not available to the agent (in virtue of the fact that they are unknown) pose no threat to conjecture two.

When it comes to personal reasons, the case for a knowledge requirement on factive emotional states is strong. The result of combining a personal reason attribution with knowledge

---

question of which of the P's that serve as personal reasons for emotional V-ing are suitable propositional complements for V-ing that. I will not present my own ideas about this issue in this paper. (Thanks to Timothy Williamson for discussion here.)

<sup>11</sup> One reader suggested that something like conjecture two is also plausible for the second type of personal reasons mentioned above. According to the reader, the fact that 'Her reason for leaving was to avoid the fire' seems to require that there was a fire. But the reader is misled by a poor choice of example. In their example, the infinitival of choice contains 'the' and 'the' carries the presupposition of a fire. (And in any case the analogue of factivity would have to do with whether the fire was avoided not whether it existed.) There is nothing akin to a general factivity requirement on these infinitival reason constructions. Consider the following example: 'Her reason for getting on that train, which sadly crashed, was to visit her relatives in New York.'

denial is as linguistically bizarre as the result of combining a factive emotional state ascription with knowledge denial. It would be strange to say, ‘Jason is angry that he is the shortest man in the room but Jason doesn’t know that he is the shortest man in the room’. And it would be similarly strange to say, ‘Jason’s reason for being angry is that he is the shortest man in the room but Jason doesn’t know that he is shortest man in the room’. Thus, we should expect that the explanation of the knowledge requirement on factive emotional states will carry over nicely to personal reason constructions. In the earlier case where Maria sees Sarah’s identical twin fall off a cliff, it is odd to say that Maria’s reason for being delighted is that Sarah fell off a cliff. The mere observation that personal reasons (as expressed by that-clause constructions) are factive does not explain the oddity, since in the cases described, it is true that Sarah fell off a cliff. A justified true belief requirement on personal reasons fares no better since Maria’s belief that Sarah fell off a cliff is true and justified. But a knowledge requirement on personal reasons provides a satisfying explanation for the oddity.<sup>12</sup>

Conjectures one and two thus seem true and when combined, offer a nice explanation of the knowledge requirement that they entail. In the next section, we shall see that while a knowledge requirement on factive perceptual states is also very plausible, the source of the requirement is importantly different.

### **2.3 Emotional Factives and Perceptual Factives**

With Unger and Williamson, I take it that perceptual verbs such as ‘seeing that’ and ‘perceiving that’ are not merely factive but are also knowledge entailing. I will not offer additional arguments for this view here. I will merely note that the case for the knowledge

---

<sup>12</sup> The idea that there is a knowledge requirement on personal reasons is not new. One will find helpful discussions of this requirement in Unger (1975), Hyman (1999, 2011, 2015), Williamson (forthcoming), Hornsby (2007, 2008), Alvarez (2010), Littlejohn (2012) and McDowell (2013), and Hawthorne and Magidor (forthcoming).

requirement on factive emotional states carries over very nicely to factive perceptual states. In the case provided in Sect. 2 where Maria sees Sarah's identical twin fall off a cliff and forms a justified true belief (via vision) that Sarah has fallen off a cliff, it would be wrong to say that Maria saw that Sarah fell off a cliff. On one level of abstraction, seeing that P and being angry that P relate to knowledge in the same way. Maria cannot see that Sarah is falling off a cliff unless Maria knows that Sarah is falling off a cliff. Similarly, Maria cannot be amused that she runs faster than John unless Maria knows that she runs faster than John. Nevertheless, I will argue that on close inspection, there is a marked difference in how emotional factives and perceptual factives relate to knowledge.

Let us begin by considering a vision-theoretic use of 'sees that'. I say, 'vision theoretic use' because there is another use of 'sees' which has nothing to do with vision.<sup>13</sup> There are two natural settings where we might say that Bukowski sees that his horse is a failure. In one setting, Bukowski sees through vision his horse lose the race and thus comes to know that his horse is a failure via vision. In another setting, we can tell a story where Bukowski learns that his horse is a failure by some means other than vision (perhaps he learns this fact by testimony). If we ever say that Bukowski sees that his horse is a failure in the second setting, then we seem to utilize a metaphorically extended use of 'sees' that communicates knowledge and little more. (Note that there are similar uses of 'sees' in connection to mathematical insight that do not involve vision). It is the first use connected to vision that concerns me here. (Similarly, there is a use of 'perceives'—in mathematics, for example—that has nothing to do with sensory perception. I will be interested in the straightforwardly perceptual use here.)

---

<sup>13</sup> For an extensive discussion of these two uses, see French (2012).

I have said that both ‘S sees that P’ and ‘S is angry that P’ entail that S knows P.

However, the analogue for conjecture one (Sect. 2) for ‘sees’ and ‘perceives’ is hopeless. In a case where Bukowski is angry that his horse has lost, one of his reasons for being angry is that his horse has lost. But in the case where Bukowski sees that his horse has lost the race, it would be wrong to say that one of Bukowski’s reasons for seeing is that his horse has lost the race.<sup>14</sup> (Relatedly, the cause-effect relation between knowledge and the relevant mental state is very different in each case. When S is angry that P, S’s knowledge that P is one of the causes of S’s anger. But in a case where S sees that P, it is not that S’s knowledge that P causes S’s seeing. It is S’s use of vision which causes S to come to know P.) Thus ‘perceives that P’ and ‘sees that P’ are not constructions that are used to convey personal reasons for being in certain states. Rather, they are devices for marking how an individual comes to know that P. In its visual use, ‘S sees that P’ communicates an episode of knowing that P via vision and in its sensory-perceptual use, ‘S perceives that P’ communicates an episode of knowing that P via sensory perception. In short, ‘S sees that P’ and ‘S perceives that P’ carry information about how knowledge was acquired whereas ‘S is angry that P’ and ‘S is happy that P’ conveys an upshot of the knowledge. Knowing P makes P available to an agent to use as a personal reason. Factive emotional state ascriptions convey what emotional state occurred for what reason(s).

In sum, factive emotional verbs like ‘anger’ or ‘regret’ require knowledge via a knowledge requirement on personal reasons but ‘sees-that’ requires knowledge because it specifies the mode of coming to know.

---

<sup>14</sup> The contrast in the text is reinforced when one considers adverbs that are connected with rational evaluation. Consider ‘foolishly.’ The felicity of this adverb patterns with states that are amenable to personal reasons: ‘He foolishly left.’ ‘He foolishly believed that she was innocent.’ ‘He foolishly regretted leaving that job.’ But ‘She foolishly saw that there was a tiger in front of her’ sounds terrible. (All of these examples pattern similarly when one replaces ‘foolishly’ with ‘wisely.’ Thanks to Timothy Williamson for discussion here.)

## 2.4 The Perceptual Model of Emotions

According to the so-called ‘perceptual model,’ an emotional experience can be evidence of value in the way that a perceptual experience of a something as looking red can be evidence of the presence of a thing that is red. As Michael Brady (2013) points out, the key issue for the perceptual model is whether or not there really is a deep analogy between the epistemological roles of emotion and those of perceptual experience. And here, we are looking for a deeper similarity beyond one obvious point of analogy that we can all agree upon (one that is not particularly epistemological): both emotions and perceptual experiences have phenomenal properties (a ‘what it’s like’ aspect). (Brady 2013, p. 47).<sup>15</sup>

One central issue for the perceptual model is whether it is right to think of emotions as representing evaluative states of affairs (as opposed to merely being prone to prompt evaluative thoughts). If they do not so represent, then it seems obvious that the perceptual model is at least, *prima facie*, in trouble. But I will not pursue this issue here.<sup>16</sup> Rather, I will press another central challenge to that model which is closely connected to the discussion thus far (though, I will now generalize beyond the case of factive emotional states): emotions and perceptual experiences seem to bear very different relations to personal reasons.

Consider a case where something looks to be red. There may be any number of explanatory reasons why the thing looks red. Some of these explanatory reasons are mundane.

---

<sup>15</sup> In passing, I should note that there is a venerable tradition—associated with James (1884) and Lange (1885)—that accepts that emotions have representational content but which holds that they represent oneself as being in a certain bodily state. On this picture, fear is a representation of oneself as being in a certain bodily state, perhaps under a distinctive mode of presentation. This kind of model hardly gives comfort to the idea that emotional states represent evaluative states of affairs. Instead, it offers a very different way of developing the idea that emotions are somehow perceptual which would be very different to the ‘perceptual model’ that is discussed in recent work on emotions. It would be interesting to explore to what extent the remarks that follow make trouble for James-Langian perceptual ideas. (Thanks to Jeremy Goodman for discussion here.)

<sup>16</sup> I will address these representational issues in chapter five.

For example, perhaps the thing looks to be red because the lighting is normal or because one's eyes are open and looking in the direction of the red thing. Some of the reasons that explain why the thing looks to be red are more sophisticated and would be revealed by psychophysics (e.g. that the wavelength of the light reflected by the object has a certain profile). But while there are plenty of reasons that explain the fact that something looks to be red, it is important to notice that an agent does not have personal reasons for perceptually representing a thing as red. In fact, the question, 'What is the agent's reason for it looking that way?' sounds fundamentally confused. This is a point that extends to other perceptual modalities. (The questions, 'What is the agent's reason for feeling warm?' or 'What is the agent's reason for the music sounding loud?' seem equally confused.)

A few points of clarification. First, note that in claiming that agents lack personal reasons for perceptual experiences, I am not simply relying on the point from earlier sections—that factive perceptual state reports do not encode personal reasons for perception. After all, there is no logical incoherence in the thought that we have personal reasons for perceptual experiences but those reasons are encoded by other means. That said, even when we widen our gaze, it is not natural at all to talk about an agent as having personal reasons for a perceptual experience. That is the point I am relying on here.

Second, I do not deny that one could have a personal reason for making something look red. But in this kind of case, one has a personal reason for an action, not for the experiential state itself.

Third, I do not deny that something's looking red might be one's reason for feeling confident that it is red. But in this case, the experiential state (or a fact about it) is the reason for one's confidence. It is one thing for perceptual experience (or facts about it) to serve as an

agent's personal reason for various other states of mind.<sup>17</sup> It is quite another thing for an agent to have personal reasons for the perceptual experience.<sup>18</sup>

We may not have personal reasons for perceptual experiences but emotional states seem eminently fit for having personal reasons. Indeed, we have seen that factive emotional state ascriptions encode the personal reason(s) for being in some particular emotional state.<sup>19</sup> John cannot ask Maria (in a natural way) what her reason was for a thing looking purple to her but it would be quite natural if John asked Maria what her reasons were for feeling angry or feeling afraid. This consideration alone seems to point to a significant disanalogy between perceptual states and emotional states.

Of course, in saying that emotional states are the sorts of things for which we have personal reasons, I do not mean to suggest that we have direct voluntary control over them. I

---

<sup>17</sup> Another issue worth exploring is the relation of emotions and experience to so-called, 'normative reasons' paradigmatically encoded by linking 'reason' with infinitival expressions (as in, 'She has reason to be angry'). Here the contrast is a little less sharp. Suppose there is an evolutionary advantage accruing to a thing's looking a certain way. Then there is, in a natural sense, a reason for it to look that way. The main point here is that normative reasons need not encode things that are even available as personal reasons. (One familiar point to make is that if a glass contains poison and I have no hope of knowing that, then there is still a perfectly natural sense in which we would say that I have reason to refrain from drinking it. This is the so-called 'external' use of normative reasons talk. While 'reason to' constructions do not straightforwardly encode the availability of personal reasons, other normative reason-theoretic talk may be more tightly connected. For example, I am sympathetic to Unger's (1975) claim that there is a fairly tight connection between talk of a state's being 'reasonable' and having personal reasons for it. In particular, it seems that a state can be reasonable only if one has personal reasons for being in it.

<sup>18</sup> Arguably, the key insight underlying historically pervasive talk of 'perception as passive' is that perceptual experiences are not the kinds of states for which we have personal reasons. Now certainly, some historically popular doctrines concerning passivity were overdone. For example, consider *Rules for the Direction of Mind*, where Descartes wrote, 'sense-perception occurs in the same way in which wax takes on an impression from a seal'. See Descartes et al. (1984, p.40) As many philosophers and cognitive scientists since have pointed out, the contours of visual experience can be affected by various cognitive factors (including priming, conceptual repertoire and so on), and so this take on passivity is misleading. But the view that perceptual experience is in a way passive because it is devoid of personal reasons may yet be defensible. (Thanks to John Hawthorne for discussion here.).

<sup>19</sup> Of course, we can also have personal reasons for emotional states that are not picked out by factive emotional verbs. For example, the fact that someone has a terminal illness might be their reason for feeling afraid. But since a proposition can be one's personal reason only if it is true, constructions like 'being afraid that P' or 'being worried that P' are ill fit to encode personal reasons, since these constructions are not factive.

might not have direct control over my feeling of shame at being seen naked by the chimney cleaner. But my feeling of shame is still something that I have a personal reason for. (Namely, my reason for feeling ashamed is that the chimney cleaner saw me naked.) This is an important way in which we conceive of emotions as responses to what one learns about the world. The attitude of belief also illustrates the fact that being governed by personal reasons is not the same thing as having direct voluntary control. It is often noted that it seems we lack direct voluntary control over our beliefs. I cannot, simply by an act of will, come to believe that the chimney cleaner saw me naked. Nevertheless, something can be my reason for believing something else. The presence or absence of direct voluntary control does not establish whether something is the sort of thing for which one has personal reasons.

One further point of clarification: In saying that emotions are the kinds of state for which we can have personal reasons, I do not want to assume that we always have personal reasons for our emotional states. By analogy, consider the case of belief. From the fact that belief is a kind of state for which it is possible to have personal reasons, it hardly follows that we always have personal reasons for belief. Turning to the case of emotions, the above discussion indicates that cases of anger that are properly describable using expressions of the form, 'S is angry that P' will inevitably encode personal reasons. But I make no assumption, and indeed, I do not think that every case of anger is so describable. (In many cases where S is angry because S falsely thinks that P, no construction of the form, 'S is angry that P' will be apt.)

While the way we talk about personal reasons is a trouble spot for the so-called perceptual model, I do not think that the preceding considerations kill all hope for that model. In this connection, let me stress two points.

First, nothing in the above discussion precludes emotions from serving as evidence for evaluative states of affairs. A very standard rough and ready account of evidence is in terms of probability raising. A fact is evidence for P just in case it raises the probability of P. Or perhaps more carefully, a fact is evidence for a particular agent that P just in case it makes P more likely from the point of view of that agent. And the main point to bear in mind here is that there is no simple argument from the premise that emotions are the sorts of things for which we have personal reasons to the conclusion that emotions cannot be evidence for evaluative states of affairs. The evidential scope of emotions in this connection needs separate consideration. And if the main idea of the ‘perceptual model’ is merely that emotions are systematically capable of playing an evidential role, the jury is still out (as far as this discussion is concerned).

Second, the above discussion has relied on observations about the ways we ordinarily talk about personal reasons. But there still remains the possibility that psychological talk for ordinary-folk is confused. Perhaps, even though we ordinarily talk as if people can have personal reasons for emotions but not perceptual experiences, there is, in fact, no psychological joint in nature that corresponds to this conceptual contrast. The question of whether our ordinary ways of talking about our psychology is apt or confused is a separate question and one which I will not worry about here.

## **2.5 Concluding Remarks**

There are striking similarities between ‘S regrets that P’ and ‘S sees that P’. Both constructions are factive and both entail that the proposition picked out by the that-clause is known. There is also a striking dissimilarity. Like other factive emotional state ascriptions, ‘S regrets that P’ encodes a personal reason for being in that particular emotional state. But like other factive perceptual state ascriptions, ‘S sees that P’ does not encode a personal reason for

being in that state and instead merely encodes a way of coming to know. The contrasting relation to personal reasons appears to extend more generally to emotions and perceptual experiences.

While emotions are the sorts of things for which we often have personal reasons, it seems that we never have personal reasons for perceptual representations. A proponent of the perceptual model of emotions will need to grapple with this fact, showing either that the perceptual model can survive this contrast or else do work to display the apparent contrast as misleading or illusory.

## Chapter Three

## Emotions, Evidence, and Safety

This discussion concerns the question of whether (and to what extent) the emotions can play an epistemically valuable role in the formation of evaluative beliefs.<sup>1</sup> In part one, I present a strong *prima facie* case for the idea that emotions can and do play an evidential role in the formation of evaluative beliefs (and relatedly that they can and do serve as good reasons for evaluative beliefs). In part two, I address objections raised by Michael Brady (2013) in his recent and extensive critique of this evidential thesis.<sup>2</sup> I argue that his objections fail to undermine the *prima facie* case for this evidential thesis. In part three, I explore some ways that the emotions can be of epistemic value without playing an evidential role. In this connection, I first outline the safety conception of knowledge and then, drawing on empirical literature, point to ways that emotions can be safety enhancers.

### 3.1 Emotions as Evidence

There is a strong case for the thesis that emotions can be evidence for evaluative states of affairs and relatedly that emotions can be *good* personal reasons for beliefs about evaluative states of affairs. Following Grice (2001), I use the term ‘personal reason’ to mark the kind of reasons that are picked out by the possessive construction (e.g. ‘my/his/her reason for  $\phi$ -ing is

---

<sup>1</sup> That is not to say they cannot be of epistemic value with respect to beliefs about non-evaluative matters, including, most obviously, beliefs about emotions. Much of the discussion to come can be extended to questions about whether and how emotions can be evidentially valuable with respect to various non-evaluative categories of belief. Much of the relevant literature has focused on the evaluative case and I follow suit.

<sup>2</sup> Brady’s (2013) skepticism about the evidential roles emotions can play *vis à vis* evaluative beliefs is used as a springboard for his extensive critique of ‘the perceptual model of emotion’. I will not address the perceptual model here as I explore these issues in chapter two.

that P') and 'explanatory reasons' to mark those reasons which are typically picked out by 'reason-why' constructions (e.g. 'a reason why S  $\phi$ -ed is that P').<sup>3</sup>

Let me begin by making a few preliminary remarks about my framework for thinking about personal reasons and evidence. (I am hopeful that much of what I have to say can be adapted to other frameworks). First, I take it that in order for a proposition to serve as an agent's personal reason, the proposition must be known.<sup>4</sup> If John's reason for leaving was that the house was on fire, it follows that John knew that the house was on fire. As a special case, a proposition about an emotional state must be known if it is to function as one's personal reason. If Robert's reason for thinking that he is a bad person is that he feels guilty that he kissed his friend's wife, then Robert knows that he feels guilty that he kissed his friend's wife. Of course, this is an observation about personal reasons, not explanatory reasons. For even if a fact about one's emotional state is not known, it can explain why other facts obtain. For example, even if John does not know about his fear, the fact that he is afraid may nevertheless be a reason that explains why John has an elevated heart rate.

I think something along these lines is also true about evidence. For a fact to serve as evidence for an agent, it is necessary that the agent knows that fact. (This is one direction of Williamson's E=K equation).<sup>5</sup> For example, John must know that he is afraid for the fact (that he

---

<sup>3</sup> While personal reasons are always explanatory, explanatory reasons are not always available for an agent to use as a personal reason.

<sup>4</sup> See Unger (1975, p. 206-14), Hyman (1999) Williamson (2000) Grice (2001) Alvarez (2010) Littlejohn (2012) and Hawthorne and Magidor (*forthcoming*). See Dancy (2000) for an alternative framework. This is not the place to argue for my preferred conception of personal reasons.

<sup>5</sup> See Williamson (2000, p. 184-208).

is afraid) to function as part of John's evidence.<sup>6</sup> Williamson provides a wealth of considerations in favor of this thesis and I will not rehearse them here.

What is it for a proposition to be evidence in favor of a second proposition? A very standard account of evidence can be illustrated in terms of probability raising. A known fact (P) is evidence for Q just in case it raises the probability of Q. Or more carefully, if an agent knows P, that is evidence for that agent that Q just in case P makes Q more likely from the point of view of the agent.<sup>7</sup> I will make use of this simple model in what follows. (This simple probabilistic model corresponds to the standard probabilistic notion of confirmation in philosophy of science. See, for example, Maher (2005). Arguably, this account of the evidence-for relation could be refined further but such refinements will not matter much to what I have to say here.)

With this framework in place, it is easy to generate a strong *prima facie* case for the idea that emotions can play an evidential role with respect to evaluative propositions. I will refer to this idea as *the evidential thesis* hereinafter.

From the outset, it should seem obvious that facts about emotions can play a probability-raising role *vis à vis* evaluative facts. Suppose Lonnie is unsure about whether he is brave and he has not accumulated very much evidence either way. When Lonnie encounters a minor danger and feels altogether petrified, Lonnie takes this as pretty good evidence that he is not brave (and he seems right to do so.)

It is easy to generate cases where facts about an agent's emotional state play a probability-raising role *vis à vis* evaluative facts about the agent. But what about evaluative

---

<sup>6</sup> Here, I am thinking of the primary use of the ideology of evidence. There is a secondary use, as in, 'the bloody knife is my evidence'. In this case, a knowledge requirement is not directly applicable because knives are not propositions. I take it that this use is derivative from facts involving the primary use.

<sup>7</sup> One will find a more detailed explication and defense of this simple model in Williamson (2000, 186-194).

propositions that concern the world external to the agent? Can facts about an agent's emotions raise the probabilities (for the agent) that certain worldly evaluative propositions are true? It certainly seems so. The fact that someone else (other than oneself) has certain emotional experiences can definitely raise the probability of certain worldly evaluative facts. Suppose I learn that the person next to me feels afraid. Learning this fact about their fear may be very good evidence that they are in danger. Given this, it is difficult to see why the fact that I feel fear cannot similarly be evidence (for me) that I am in danger. Similarly, the fact that I always feel amused when I am in the presence of Jones is pretty good evidence that Jones is amusing. It thus appears that there is a strong case for the idea that emotional facts can play the probability-raising role that we associate with evidence.

We can make the case more abstractly. Suppose S knows of a strong positive correlation between one phenomena (P1) and another phenomena (P2) and then learns that P1 obtains. Unless S already knows that P2 obtains, learning that P1 obtains tends to raise the epistemic likelihood for S that P2 obtains. Now let P1 be a certain kind of emotional fact, and P2 a certain kind of evaluative fact. Those who deny that the former is evidence of the latter will be forced to say one of three things:

- (i) We never know about a strong positive correlation between the emotional state and the evaluative state.
- (ii) We never know when the relevant emotional fact obtains.
- (iii) Whenever we do know that the relevant emotional fact obtains, we already know about the evaluative state or we have evidence that screens off the evidential value of the emotional state.

The strong *prima facie* case for the evidential thesis is that we can generate cases where (i), (ii) and (iii) are not plausible options. In the case of 'fear', (i) should not be taken seriously. As for (ii), sometimes we know about our fear and so (ii) is a non-starter. And given the fact that

we can suddenly come to know that we are afraid without already knowing that we are in danger (or without having evidence that makes the fear evidentially worthless), (iii) is also implausible.<sup>8</sup>

The *prima facie* case above points towards another question worth asking: Are emotional facts ever *good* personal reasons for believing that some evaluative fact/state of affairs obtains?<sup>9</sup>

As a crude first pass, one might think that insofar as P is evidence for Q, it is appropriate to use P as one's reason for believing Q. But this idea requires some sharpening. For P might raise the probability of Q from .1 to .2 but that would not make it appropriate at all to believe Q.<sup>10</sup>

However, insofar as emotional states can play a probability raising role with respect to evaluative propositions, it is difficult to see why there would not also be occasions where the probability of an evaluative proposition (whose probability is raised) is already reasonably high. In such a case (where P raises the probability of Q, and Q has a high resultant probability), it seems right to say that P is one of the *good* personal reasons that the agent has for believing Q. (This is especially clear in situations where the agent appreciates the probabilistic import of P. Matters are a little more delicate in cases where P makes Q more likely from the agent's perspective but the agent does not realize it.)<sup>11</sup>

---

<sup>8</sup> It is worth mentioning that it may be a bit naive to suppose that there is a strong positive correlation between fear and one's being in danger. For one might feel fear because of a danger to someone one cares about. Of course, it is still plausible that one can come to know that there is a strong positive correlation between certain kinds of fear and certain kinds of dangers. Similar points arise for correlations connecting other emotions to other evaluative states of affairs.

<sup>9</sup> Here, I use the expression 'good personal reasons' on account of the fact that some personal reasons for belief/action are simply bad reasons to believe/act.

<sup>10</sup> Suppose John is told that sandwiches will be served at lunch. This fact might raise the epistemic probability that exotic fish sandwiches will be served at lunch from .1 to .2 but the fact that sandwiches will be served at lunch is not the kind of fact that, by itself, makes it appropriate to believe that exotic fish sandwiches will be served at lunch.

<sup>11</sup> There is another category of reasons worth mentioning; Normative reasons are typically expressed by the 'reason-to' construction. These are different from what I call, 'good personal reasons' since one might have a good reason to believe P and yet not believe P and thus, have no personal reasons for belief at all. We can easily adapt the above discussion to argue that emotional facts can be good normative reasons to believe various evaluative propositions. For example, if one knows P and P raises the probability that Q,

### 3.2 Michael Brady's Objections

I have presented the *prima facie* case for the evidential thesis (and relatedly, for the thesis that emotions can serve as good personal reasons for belief). Let us turn to Michael Brady's (2013) case against the thesis that emotions can serve as 'bona fide reasons' for evaluative beliefs (and more generally, his case against the idea that emotions can play an evidential role in the formation of evaluative beliefs).<sup>12</sup> Brady's argument can be usefully distilled into four main objections. I will present his objections, and reply to each in turn.

#### The Conclusive Reasons Objection

Let us begin by considering, what I will call, Brady's 'conclusive reasons objection'. Here, his main observation is that unlike emotional experiences, people are inclined to treat their perceptual experiences as conclusive reasons/evidence for the corresponding perceptual beliefs. In other words, when an agent has no defeating evidence, perceptual evidence suffices as conclusive justification for belief. For example, if it appears to S that there is a coffee cup on the table, then barring defeaters, the appearance alone fully justifies S's belief that there is a coffee cup on the table. According to Brady, we do not regard our emotional experiences as conclusive evidence in the same way.

There are a few obvious responses to this objection. First, even if this observation is true, it does not undermine the modest epistemological thesis that some facts about emotions can be evidence for evaluative propositions. For even if emotions can provide evidence for evaluative beliefs, it is a further question whether they provide sufficient evidence to fully justify a belief.

---

and the resultant probability of Q is high, then the proposition that P is a good normative reason to believe Q.

<sup>12</sup> Brady does not explain what he means by 'bona fide' reasons. Because Brady's ideas about the nature of reasons and evidence is somewhat underdeveloped in the text, I am considering his ideas through the lens of my framework for thinking about reasons and evidence.

Let us say that evidence for P is strong when the evidence makes it likely that P and evidence is weak when it merely raises the probability that P. The fact that Frank is over five feet tall is strong evidence that he is taller than an average puppy. And the fact that Frank lives in New York is weak evidence that he is a cool guy. In general, weak evidence is not the sort of thing that, by itself, makes it reasonable to believe P (whereas strong evidence does, by itself, make it reasonable to believe P). The evidential thesis merely requires that emotions can serve as weak evidence for various evaluative propositions.

Second, it is worth noting that weak evidence can be combined with additional evidence to generate high epistemic probability; so, weak evidence can still function as a good personal reason for believing P if it is in the company of other evidence that makes it likely that P.

Third, we should be doubtful about whether Brady's main observation is true in full generality. If Lonnie learns that he is terrified at the thought of holding a bunny, it would be reasonable for Lonnie to treat that fact as pretty conclusive evidence that he is not brave. Similarly, if I feel repulsed at the thought of kissing Frank, I might very well take that as conclusive evidence that I should avoid Frank's kiss. Ordinary people do not take such emotions as preliminary gambits to then be replaced by sophisticated arguments or further inquiry.

### **The Motivating Response Objection**

Let us turn to a second reason why Brady thinks we do not treat our emotional experiences as sources of justification. Brady thinks of the emotions as 'automatic and reflexive responses' that, unlike perception, 'capture and consume our attention' in a way that motivates us to seek and discover further considerations that will bear on their accuracy. Brady writes,

Consider the experience of fear, when trying to get to sleep at night, upon hearing a noise downstairs. In such circumstances we are motivated to seek out and discover additional reasons or evidence. In particular, we are motivated to seek out and discover considerations that have a bearing on whether or not our initial

emotional appraisal-namely that we are in danger- is accurate. We strain our ears to hear other anomalous noises, or rack our brains trying to think of possible non-threatening causes for the noise, and so forth. It is unlikely that, in these circumstances, we would regard our feeling of fear as conclusive reason to judge that we are in danger. (Brady, 2013, p. 139)<sup>13</sup>

According to Brady, emotions merely facilitate the search for reasons/evidence so that we may assess the accuracy and rationality of our emotional responses. Thus, emotions do not confer justification for their corresponding evaluative beliefs. Instead, emotions are motivational resources that direct and focus our attention towards various non-evaluative features of our experience (e.g. the sharp teeth of the large dog that is leaping towards you), and this focus allows us to acquire evidence for our evaluative beliefs. For Brady, emotions do not justify beliefs; emotions merely motivate the search for justification.<sup>14</sup>

Brady is right to think of emotions as automatic and reflexive responses that capture and consume our attention, and direct us to the most relevant stimuli in our environments. (At least in cases where the emotional state is caused by the things in our environment. Brady rarely discusses those cases where we have anticipatory emotions that are not prompted by environmental cues.) Yet, Brady's emphasis on the motivational role of the emotions fails to show us how/why the emotions cannot also play a justificatory role with respect evaluative beliefs. It would be a *non-sequitur* to argue from the fact that emotions play one kind of epistemic role (motivating the search for further reasons/evidence) to the conclusion that emotions do not play any other epistemic roles (evidential or not). Brady also overestimates the contrast with perception here. Visual perception can (and often does) motivate the search for

---

<sup>13</sup> A longer version of this example occurs in chapter three (2013).

<sup>14</sup> Brady offers an alternative way that the emotions may play a positive epistemic role- mainly by being a motivational mechanism which promotes what he calls 'evaluative understanding.' (2013, Ch. 4.) I will not consider this alternative account here.

further evidence too. If it appears that something is about to fall over, one might seek out tactile reassurance of its sturdiness; if it appears that someone does not like you, one might immediately seek out confirmation of one's pessimism; if something appears to be a large diamond, one might seek out confirmation that it is not a fake. Meanwhile, as noted, emotions do not always prompt the search for further confirmation (consider the feeling of revulsion at the thought of Frank's kiss described above).

Further, if some emotional state motivated a search for further reasons to believe P, then that positively suggests that the emotional state provided some initial reason to believe P. Brady voices the intuition that emotions are usually taken to be unreliable and inconclusive phenomenal experiences that, at best, direct and focus our attention to find reliable and conclusive reasons.<sup>15</sup> Yet, the idea that an emotional state could motivate the search for additional reasons to believe P seems to support the idea that the emotional state gave us a reason for believing P in the first place.<sup>16</sup>

It seems that Brady is faced with a dilemma. If emotions cannot be reasons for believing, then we are not rational in allowing our emotions to prompt any kind of inquiry at all. Recall Brady's case of someone trying to go to sleep and then suddenly feeling fear upon hearing a noise downstairs. Brady says our emotions motivate us to seek additional evidence that bears on our initial emotional appraisal. Later, Brady writes,

---

<sup>15</sup> Such reasons enable us to assess the accuracy and rationality of the emotional response so we may render the corresponding belief or judgment as warranted.

<sup>16</sup> Elgin (2008) argues that some emotions can be reliable indicators of certain aspects of their objects. The experience of fear can be evidence enough to believe one is in danger, or the feeling of trust can be evidence that one finds something to be reliable. It is plausible that the very fact that emotions present themselves as indicators at all means we sometimes afford them an initial degree of epistemic access to their objects. Brady is aware of Elgin's initial tenability thesis and answers it with his Double Counting argument.

We might think instead that it is normal for us to endorse the content of our emotional experience only after we have discovered (or on some occasions invented or fabricated) reasons to think that our emotional appearance is veridical. (Brady 2013, p. 142)

If our emotional experiences are as indiscriminate as Brady seems to believe, then we should never trust our emotional experiences in the first place (let alone, allow our emotions to be proxies that motivate any action whatsoever). After all, on Brady's account, emotions are not themselves reasons for believing anything (they are merely responses to features). Given Brady's case, it seems one would be more rational to just ignore their fear and go back to sleep. Brady responded to this objection in correspondence:

The emotion might raise a possibility, or even raise the probability, without being evidence, I think. So fear when waking up at night might raise the possibility that I'm in danger; but the fact that something raises a possibility isn't a reason to think the possibility likely, and so a reason to believe the relevant proposition. Consider: the fact that the barista is genuinely pleased to see me raises the possibility that she is the love of my life. But it's surely not evidence for this, or a good reason to think that she's the love of my life. It is rational, I think, to search for reasons and evidence raised by even slight possibilities: knowing that there's a slim chance I'll be bitten by a shark, it is rational for me to consult the lifeguard about shark-infestation in these waters. But again the fact that it's possible that I'll be attacked by a shark isn't a reason to believe that I will be.<sup>17</sup>

There is a substantial difference between raising the possibility that P and raising the probability that P. Raising the possibility that I get sick is a matter of making that possibility salient. Contrast this with the notion of raising the probability. One application of this ideology involves objective probability, where something happens in the world that makes something else more likely to happen (e.g. the fact that John ate rotten food makes it more likely that John will get sick.) This is objective probability raising. But in this context, the epistemic use of probability is more relevant: An agent acquires new information such that the epistemic likelihood of P (relative to this new information that P) is now higher than it was relative to the

---

<sup>17</sup> Email correspondence on November 27, 2012.

agent's prior information. That seems to be what Brady is alluding to here. But a new piece of information is, at least, weak evidence for P when it raises the evidential probability in this way.

For the most part, Brady seems to distance himself from talk of probability raising. Brady hopes to convince us that emotions are merely possibility-raisers. That is, an emotion like fear merely makes *the possibility that one is in danger* a salient one, thereby putting one in a position to investigate their situation further. But these investigations seem best explained by supposing that emotions have some evidential value.

### **The Proxy Reasons Objection**

Brady presses another point that I will call his 'proxy reasons objection'. Insofar as emotions are reliably correlated with certain evaluative features, it is because they reliably track certain non-evaluative features that make for those evaluative features. But in such a case, an agent's reasons for believing that some evaluative state of affairs obtains is not given by the emotional state but by the agent's awareness of the features that their emotional state tracks. Given this, Brady thinks of emotions as 'proxy' or '*pro tempore*' reasons.

To put things slightly differently: when we believe on the basis of emotional experience, we are believing on the basis of *pro tempore* reasons, reasons "for the time being", which we rely upon precisely because we presently lack awareness of features that constitute genuine reasons for our judgments. For note that if we did become aware of such features, then the justificatory force of our emotional experience would seem to disappear. This is because the justificatory force of our emotional experience would be exhausted by the justificatory force of the features that we take our emotional experience to reliably track. It is, to repeat, only because our emotions reliably track such features that they can function as signs or indicators of value. (Brady 2013, p. 130)

This argument involves two mistakes. First, I suspect that Brady somewhat conflates explanatory reasons and personal reasons. We should not have any expectation that a reason for believing that a state of affairs obtains should also be an explanatory reason why that state of affairs obtains. If John looks to be in pain, that is evidence that John is in pain and a reason to

believe that he is in pain. But the fact that John looks to be in pain is not one of the explanatory reasons why he's in pain. Similarly, the fact that someone tells me that I am in danger is a reason to think that I am in danger but it is not an explanatory reason why I am in danger. In general, the fact that P is not an explanatory reason why Q does not prevent P from being an agent's personal reason for believing Q (and does not prevent P from being evidence for Q). Although the text is not decisive, I suspect that Brady's motivation for demoting emotions to being mere 'proxy reasons' is that they will not figure in the proper explanation of why things are a certain evaluative way. But as we have seen, if this is his motivation, it is a bad one.

The second mistake is more straightforwardly visible in the text above. Here, the mistake is in thinking that emotions cannot be personal reasons on account of the fact that if we became directly apprised of the features that emotions 'track', the evidential force of the emotions would disappear. The main point I hope to make is that even if the evidential force of the emotion disappears (upon being presented with certain facts tracked by the emotions), that does not show that the emotions were not, in the initial state, genuine reasons/evidence for belief.

An analogy with testimony is helpful here. Suppose a reliable person tells me that I am in danger. By any reasonable standard, this testimony is evidence that I am in danger and is a reason for me to believe that I am in danger. Now suppose what the testifier is sensitive to is that there is a hungry tiger behind me. It is certainly true that if I were aware of this fact then the evidential value of the testimony would disappear. But in the case where I am not aware of the fact that there is a hungry tiger next to me, the testimony is pretty good evidence that I am in danger and a pretty good reason to believe that I am in danger. (Similarly, if someone tells me that there is a cup on the table, the testimony is evidence but becomes evidentially worthless to me when I see that there is a cup on the table.)

## The Double Counting Argument

Finally, let us turn to what is arguably the most interesting line of thought in Brady's attack on the evidential status of emotions. Brady argues that when an emotional response is accompanied by an evaluative judgment, the reasons that explain the emotional response are typically the same reasons one has for reaching the evaluative judgment, namely certain perceptual facts. Brady writes,

This highlights the fact that considerations which constitute reasons for evaluative judgments are equally reasons for the relevant emotional response. So, reasons to think that something is dangerous are equally reasons to fear that thing; reasons to think that one has just been insulted are equally reasons to be angry; reasons to think that one has done something wrong are equally reasons to feel guilt; and similarly, for other central cases of emotions. (Brady 2013, p. 113)

Brady's worry about counting the emotional state as evidence can be put as follows. Insofar as one supposed that the emotional state provided an additional probabilistic boost, that would be double counting. Consider an example. Suppose Barry is fully aware that Robert is eating off the bathroom floor, which makes it likely from Barry's point of view that Robert is doing something disgusting. It would be strange if Barry reasoned as follows: 'Robert is eating off the bathroom floor and that is certainly a reason to think that Robert is acting in a disgusting manner. And Robert's eating off the bathroom floor is making me have a feeling of disgust. This feeling of disgust is giving me a further reason to think that Robert is acting in a disgusting manner.' We are thus encouraged to think that the emotional state does not provide any additional evidential boost. I call this Brady's 'Double Counting Argument'.

The very best setting to run a kind of double counting argument is one where an evaluative proposition figures in the content of a factive emotional state. For example, suppose

that I am angry that my husband has acted badly. In this case, it would be strange to suppose that the fact that I am angry that my husband has acted badly has probabilistically boosted the likelihood that he has acted badly. And this is because being angry that my husband has acted badly arguably requires the fact that my husband acted badly to serve as one of my personal reasons for feeling angry, which in turn requires my knowing that he acted badly.<sup>18</sup> But if I already know that he has acted badly, it is hard to see how my anger will provide me with an evidential boost with respect to the proposition that my husband acted badly. (This is especially clear on Williamson's conception of probability according to which the epistemic probability for a proposition is its likelihood on what I know. Thus, if I know P then the likelihood of P is already maximal.)<sup>19</sup>

The second-best setting to run the double counting argument is one where the evaluative proposition is not already known but the agent is knowledgeable about the perceptual cues which caused the emotional state and that knowledge clearly supports the agent's evaluation. The example above (Robert's disgusting bathroom activity) fits this format but even in this kind of setting, the double counting argument is far from secure. Consider an analogy. Suppose someone reads some graffiti and laughs. Between the act of reading and the subsequent laughing is the reaction of finding what is written funny. It does not seem right to argue that *finding a written joke funny* played no evidential role with respect to the proposition that *the joke is funny* on the grounds that the perception of the words does all of the evidential work. On the contrary, the written words count as considerations in favor of the view that the joke is funny precisely

---

<sup>18</sup> These remarks are suggestive of a more basic principle; If S is angry that P then S knows that P. This view is defended in length in Unger (1975), Williamson (2000), and Gordon (1987).

<sup>19</sup> See Williamson (2000, chapter 3). For more on the connection between factive emotional states and personal reasons, see Dietz (2017).

because one finds it funny. Once it is conceded the humorous reaction is evidentially relevant, why not also, by analogy, say that the disgust reaction is evidentially relevant too?

The above considerations suggest that even when an emotional reaction is triggered by perceptual cues and one is fully aware of those cues, it is by no means clear that the perceptual cues screen off any evidential role for the emotions. Moreover, it is worth noting that there are other types of cases that are even less vulnerable to Brady's double counting argument.

First, consider the fact that even if emotions are caused by one's perceptual engagement with the world, emotional experiences can be responses to any number of subtle perceptual cues that we are not aware of. LeDoux (1989; 2000) demonstrated that a visual stimulus could trigger an emotional response prior to perceptual awareness of salient stimuli. Since there are cases where an emotional response precedes awareness/conscious assessment of the relevant perceptual cues, it seems plausible that there are settings in which an agent's fear is the first reason the agent has for believing that the agent is in imminent danger. (This perspective is particularly appealing once one adopts the Williamsonian view that known facts can play this evidential role. If one is unaware of the perceptual cues at the time of the fear, the perceptual cues cannot serve as one's evidence.)

Second, it is worth considering the different processes by which emotion and perception occur. Emotional experiences can occur via pathways that bypass the neocortex.<sup>20</sup> Like the case described above, an emotional stimulus can be processed in a 'quick and dirty' route from the sensory thalamus to the amygdala (the 'low road'), and allow one to respond to potentially dangerous stimuli before they are perceptually aware of what they are responding to. An emotional stimulus can also reach the amygdala through a longer process (the 'high road') by

---

<sup>20</sup> The neocortex is a higher order processing system that is thought to be involved with thinking, reasoning, and consciousness. See LeDoux (1989; 2000).

way of processing information from the sensory thalamus, to the sensory cortex and then to the amygdala.<sup>21</sup>

Although the thalamic system cannot make fine distinctions, it has an important advantage over the cortical input pathway to the amygdala. That advantage is time... The thalamic pathway is thus faster. It cannot tell the amygdala exactly what is there, but it can provide a fast signal that warns that something dangerous is there (LeDoux, 1996).

If this is right, then emotional experiences are autonomous sources of justification in virtue of having an independent process by which they can interpret stimuli. If we can differentiate perception and emotion by their autonomous ways of processing sensory information, then emotion and perception seem to be independent (though closely related) sources of evidence.

While the double counting argument is the most interesting of Brady's arguments, it fails to support his claim that emotional states never serve as genuine reasons/evidence for evaluative judgments. Returning to the *prima facie* case (§1), Brady seems to take option (iii) by arguing that whenever we do know that the relevant emotional fact obtains, we already know about the evaluative state or we have evidence that screens off the evidential value of the emotional state. I hope to have shown that option (iii) is not viable across all possible cases.

### 3.3 The Emotions and Safety

I have defended the thesis that emotions can play an evidential role in connection with beliefs about evaluative states of affairs. I have also defended the companion view that emotions

---

<sup>21</sup> There is a wealth of evidence that suggests the amygdala processes early emotion signals irrespective of attention (Anderson & Phelps 2001; Todd et al. 2012; Vuilleumier and Schwarz, 2001) and irrespective of awareness (Morris et al., 1998; Whalen et al. 1998), and that the early emotional response influences a range of cognitive functions, including perception, attention, and memory. Research also suggests that there are specialized subcortical pathways that allow for the early production of emotion so that the amygdala can perceive potential threat in the environment prior to the completion of standard perceptual functions (Romanski & LeDoux, 1992). These findings strongly support the idea that the processing of emotion occurs prior to a full cognitive analysis (Zajonc, 1994).

can be *good* personal reasons for evaluative beliefs. However, it would be a mistake to suppose that this is the only positive epistemic role that emotions can play, or even their primary epistemic role. In this final section, I will explore ways that emotions can be epistemically valuable (without being evidential) within a safety-theoretic framework for knowledge. I will begin by briefly sketching out a safety framework. I will then highlight some empirical evidence for the idea that emotions can and often do enhance the safety of our beliefs.

### **Safety and Knowledge**

True belief is not sufficient for knowledge. The safety-theoretic account (popularized by Williamson (2000) but also defended in some form by Sosa (1999) and others) gives a picture of what, at least in part, is required beyond true belief. We can capture the basic idea with the following principle:

**(Crude Safety):** *S knows a true proposition P only if S could not have easily falsely believed P.*

As a number of philosophers have noted, this crude way of articulating safety needs refinement.<sup>22</sup> Crude Safety cannot account for cases where consistent with knowing P, one could have easily falsely believed P via a totally different method. Consider an example. Suppose John knows that I am in the coffee shop by calling me and asking. It could easily have happened that I was not in the coffee shop but John believed that I was on the basis of his views about my usual whereabouts. Thus, John could have easily falsely believed that I am in the coffee shop despite the fact that John knows that I am in the coffee shop (by calling and asking). Given the nearby possibility where John believes falsely, John's belief fails the necessary condition proposed by

---

<sup>22</sup> For a more extensive discussion, see Rabinowitz (2011).

Crude Safety. To preserve John's belief as having the status of knowledge, we can make a natural refinement to the Crude Safety principle by taking the method of belief-formation into account.<sup>23</sup>

Let (M) be the actual process by which one comes to believe P:

**(Refined Safety):** *S knows P via M only if S could not have easily falsely believed P using a relevantly similar method to M.*

In the coffee shop case above, the nearby false belief is produced by a very different method (relying on general beliefs about my usual whereabouts) to the method John uses in the actual world (calling and asking) and so becomes irrelevant. Refined Safety can provide a useful explanation of knowledge failure in a variety of widely discussed cases in epistemology.

*Brains in Vats.* Suppose I have a true belief that I have hands. Does the possibility of my being a brain in a vat show that my belief is not knowledge? The safety account suggests that this possibility is compatible with my belief being knowledge and this is connected to two points; First, while the brain in a vat scenario is possible, it does not follow that it is a close possibility (one that could have easily happened). Second, one might reasonably argue that methods should be externally individuated. Thus, the process (M) includes not just what goes on in the brain but what leads up to the triggering of the optic nerves, etc. Even if a brain-in-a-vat possibility is close, it is debatable whether the method by which I form the belief that I have hands is relevantly similar to the method in the brain-in-a-vat scenario.

*Fake Barns.* Suppose I have a true belief that there is a barn in front of me despite the fact that there are fake barns in the surrounding area. Here, the standard judgment is that I do not know there is a barn in front of me. The safety account can make sense of this judgment since it

---

<sup>23</sup> Thanks to John Hawthorne for extensive discussion here.

is natural to think that in the fake barn scenario, I could have easily formed a false belief that there is a barn in front of me by looking at a fake barn.

*Gettier Cases.* Suppose I look at a stopped clock at 3:00 PM and the clock reads that it is 3:00 PM and I thus come to have a justified true belief that it 3:00 PM. The standard judgment is that this is not a case of knowledge. The safety approach explains this. I could have easily formed the false belief that it is 3:00 PM by looking at the clock a little earlier or a little later, when it would have still been stopped at 3:00 PM.

*Lotteries.* Suppose I have a justified true belief that I will lose the lottery. The standard judgment is that I do not know that I will lose. The safety approach potentially explains this by appealing to the fact that while there are not many close possibilities where I win, there are some where I do win (and yet believe that I will lose using a relevantly similar method). In a lottery case, this is enough to stop my actual belief from being knowledge. Consider one further refinement that can cover even more cases.

**Super-Refined Safety:** *S knows P via M only if S could not have easily had a false belief about whether P using a relevantly similar method to M.*

Suppose John picks a true answer to a mathematical question by the toss of a coin. John is asked whether twelve multiplied by eleven equals one hundred and thirty-two and using this method, he comes to believe that it does. Clearly, John could not have easily had that belief falsely since twelve multiplied by eleven could not have easily failed to equal one hundred and thirty-two. But he could have easily had a false belief about *whether* twelve multiplied by eleven equals one hundred and thirty-two (in the scenario where he flips the coin and comes to believe the negation).

Of course, this is not a counterexample to Refined Safety since Refined Safety is only stated as a necessary condition on knowledge. But if one likes the Refined Safety principle, it seems one should also like Super-Refined Safety, and even think of it as explanatorily more powerful than Refined Safety because it can explain knowledge failure in a wider range of cases.<sup>24</sup>

It is worth highlighting two ways that the safety approach is different from most standard reliabilist approaches. First, most standard reliabilists require a high but not perfect success rate across relevant alternative cases. But safety requires a *perfect* success rate across relevant alternatives.<sup>25</sup> (For this reason, safety entails truth, since by any standard of closeness, a belief is close to itself.)

Second, this means that insofar as a safety approach is not skeptical, one's understanding of what counts as 'close' must only include a narrow set of cases. Consider a color-vision case where S observes an object from a close distance and judges that the object in front of S is red. For a case like this, let us say that a reliabilist only requires that S's belief is true across a wide range of cases at least 80% of the time (this will be our reliabilist's necessary condition on knowledge). Our reliabilist can allow that the relevant alternative cases (over which reliability is judged) includes pretty much any case in S's adult life where S judges that a thing is red (at fairly close range using color vision). But a safety theorist would be extremely hesitant to allow all such cases to count as close. For if one does, then a single false belief that something is red (in

---

<sup>24</sup> If one likes Super-Refined Safety, one might also be tempted by other conceptions of safety that can potentially offer even more explanatory power for certain kinds of cases. For example, one might craft a test that looks not merely at whether one could easily have formed a false belief about P -- using a relevant similar method -- but whether one could easily have formed a false belief about a relevantly similar subject matter. I will not pursue these issues any further here.

<sup>25</sup> One must be precise about what counts as a 'close case' if one hopes to avoid wayward alternatives that lead to skepticism. See Greco (2007) for a nice exposition of Pritchard's (2005) ambiguous formulation of 'close cases'.

one's adult life and based on color vision) would be sufficient to rule out all of the cases from having the status of knowledge. Thus, a non-skeptical safety theorist will be *much* stricter about how similar a case must be to count as a close case that deploys a relevantly similar method.

Similarly, consider a case where one forms a belief by testimony. If one allows all testimony to count as a relevant alternative or even all testimony from the same informant, then one piece of false information from that informant in the past will preclude knowledge from that informant in the future. Even a restriction to any case of testimony about the same question would usually result in pretty skeptical consequences. Falsely believing P on the basis of false testimony at any time in the past will, even if the incident was long forgotten, preclude one knowing  $\neg P$  by testimony any time in the future. Again, the safety theorist will likely have a much stricter conception of 'close cases' (for example, perhaps only cases involving very small tweaks to the recent history or current environment will count as close). Let us call this narrow conception of close cases (suggested by a combination of anti-skepticism and the requirement of being error-free across close cases where a relevantly similar method is deployed) the phenomenon of *narrow relevant alternatives*.

### **Emotions and Safe Belief-Forming Processes**

There are two structural features of the safety approach that are suggestive of a unique way that the emotions can be knowledge-conducive (without being evidential). The first is that nothing about the safety principle requires that knowledge be based on evidence. A process can be a safe pathway to belief without anything in that process being evidence for that belief. Second, and relatedly, something can play an integral role in a safe process without also playing an evidential role for the belief produced by that process. For example, suppose John has an ultra-reliable neural mechanism that generates true beliefs about which day of the week certain

dates fall on. If John does not know about that neural process or know about the correlation between that neural process and the facts about the dates, then facts about the neural process cannot feature in his evidence for believing those date-propositions. Nevertheless, the neural process plays a crucial role in the explanation of why those beliefs are safe.

These structural considerations point to a unique way that the emotions might be epistemically valuable: namely, by being an integral component of a safe belief forming process. We can illustrate this idea with an imaginary case. Suppose there is simple creature such that whenever the creature is afraid, the fear causes the creature to have a belief that it is in danger. Suppose further that there is a very reliable connection between the creature's fear and actual danger in the world such that the creature could not have easily been afraid unless the creature was truly in danger (i.e. there are no close cases where the creature feels fear but there is no danger). On any particular occasion where the creature feels fear and subsequently believes that it is in danger, the safety of the creature's belief is nicely explained by the fact that the only scenarios in which the creature could have easily had a false belief (that it was in danger) are those where the creature based its belief on something other than its feeling of fear (in which case, the creature would not have formed its belief via a relevantly similar method to the actual case).

The story above does not require (in any way) that the creature's fear provided the creature with evidence that it is in danger. For one thing, it is not part of the story that the creature is aware that it is afraid. After all, a state can causally influence one's doxastic life without one's awareness of it. And to the extent that the creature is unaware of the fact that it is afraid, it is less natural to describe the fear as evidence (especially if one adopts the view that a fact can serve as evidence only if it is known.) Even if the creature is aware of the fact that it is

afraid, the causal connection between the fear and the belief may not be underwritten by any appreciation on the part of the creature's knowing about any robust connection between fear and actual danger, in which case it is far less plausible that the fact of fear, even if known, plays a probability raising role. (More generally, we should not expect that any causal ingredient in a safe process could or would be evidence that raises the probability of the belief formed by that process.)

I have presented a sketch of how emotions might play a safety-theoretic role in the abstract but we can put empirical flesh on this idea by turning to the empirical work on the emotions. In the next section, I will sketch out various ways that the empirical literature encourages the view that emotions can and often do enhance the safety of our beliefs.

### **Empirical Support**

Let us now turn to some relevant threads of investigation in the empirical literature concerning the emotions. While the literature contains myriad themes that can be deployed to show how emotions function as safety-enhancers, I will (without being exhaustive) gesture at a few relevant threads of empirical investigation. Such data tells heavily against a blanket dismissal of the epistemic value of emotion and especially tells against blithe generalizations about passions making us more irrational.

*Enhancing Levels of Attention.* Emotions can make an epistemic difference by enhancing attention in two different ways.<sup>26</sup> First, emotions can change the depth of one's *cognitive* attention to a problem. In this connection, there is a wealth of empirical research concerning how

---

<sup>26</sup> As I indicate, not all emotions enhance attention. While fear heightens attentional focus and boosts information processing which is thought to be beneficial to the overall survival of a species (LeDoux, 2000; Dolan, 2002; Phelps, 2006), studies show how disgust diverts (sensory) attention from the target source of disgust (Liu, Zhang, et al., 2015; Krusemark & Li, 2011) (thereby serving its own evolutionary benefit by minimizing possible exposure to and threats like poison or germs). (Rozin and Fallon, 1987).

emotions can affect “depth of information processing” (Lerner, Yi, Valdesolo and Kassam, 2015; Schwarz & Bless et al., 1991.) In particular, studies (Tiedens and Linton, 2001) suggest that emotions which involve appraisals of “high-certainty” (like happiness, anger, and disgust) cause one to engage in heuristic processing (a fairly automated superficial kind of reasoning), while so-called “low-certainty” (Stalnaker, 1984) emotions (like fear or sadness) cause one to engage in deeper systematic processing. Second, emotions can also enhance one’s *perceptual* attention to particular objects or features of their environment.<sup>27</sup> For example, studies have shown that fear significantly modulates perceptual attention and focus in a way that dramatically impacts the kind of information that is encoded in memory. Some studies have focused on threatening stimuli (like weapon effects) to demonstrate how fear can enhance perceptual attention and focus (on the threat) in a way that also enhances one’s memory of specific details about the threat (i.e. what the gun looked like, the size, etc.)—when damage to the amygdala blocks one’s feeling of fear, perceptual attention, focus, and memory of the threatening stimuli is severely diminished. (See Stanley, Ferneyhough, and Phelps, 2009 for extensive discussion.)

Of course, it would be naïve to suppose that heightened attention is *always* a good thing. When one’s emotional state causes one’s cognitive or perceptual attention to narrowly focus, one may neglect other things of importance. For example, fear upon seeing a gun may enhance one’s attention and memory of specific details about the weapon but it does so at the expense impairing one’s attention to other important peripheral details of the scene (e.g. attention to the face of the person holding the weapon). (Kramer, Buckhout, and Eugenio, 1990; and Kensinger, 2009.) Also note that it would be wrong to assume that automated processes will always produce results that

---

<sup>27</sup> In passing, I note that Brady’s discussion talks as though the exclusive role of emotions is to modulate perceptual attention but as I note in the text, it is not true, in general, that every type of emotion plays this role. Moreover, it is not clear why he focuses on emotions as modulating perceptual attention when emotions can have other effects (e.g. enhancing memory, depth of information processing, etc.)

are worse than those produced by careful cognitive and perceptual attention. There are plenty of cases where enhanced cognitive attention to a question makes for a poorer job at answering it. (For discussion, see Bodenhausen, Gabriel, and Lineberger, 2000. Also think of your favorite example of a philosopher who gets things wrong by overthinking.) However, the fact that an emotion can enhance cognitive or perceptual attention to something at the peril of attending to something else (that is of relevant importance) does not preclude the emotion from facilitating knowledge.

Suppose an emotion causes one to have enhanced perceptual awareness and thus form a belief. When asking whether that belief is safe, a typical safety theorist *will not* consider the truth value of any old belief that was also caused by an emotion that enhanced perceptual attention. That coarse-grained approach is far too skeptic-friendly for the standard safety theorist's tastes. Instead, a standard safety theorist would consider whether relatively small perturbations in history and environment would produce a false belief about the relevant question using a relevantly similar method. In many cases where emotion enhances attention, it is likely that a belief will pass this kind of safety test with flying colors.

Note that when emotions facilitate knowledge through attention enhancement, they do not, in general, serve as evidence. This is true regardless of whether or not the knowledge induced is of an evaluative or non-evaluative proposition. Suppose fear prompts me to pay careful attention to whether I have enough money in the bank to afford my next three mortgage payments, and enhanced attention also causes me to make a correct calculation in a case where casual attention would have likely led me astray (I am a sloppy arithmetician). It is completely wrong to think that fear is somehow a part of my evidence for the claim that I have enough

money in the bank to afford my next three mortgage payments. (This point applies similarly to cases in which emotion facilitates knowledge of evaluative propositions by enhancing attention.)

*Enhancing Retrieval.* Emotions can also affect one's capacity to retrieve information. The basic idea here is that some information is marked/tagged in memory as important because of its emotional significance and this, in turn, makes a difference to the richness of information retrieved (Kensinger, Garroff-Eaton, and Schacter, 2006) as well as the facility and speed with which the information can be retrieved. (Richter-Levin and Akirav, 2003; Buchanan, 2007). Insofar as retrieval makes a difference to one's capacity to, say, figure something out (i.e. coming to know the answer to some question) it is clear that emotions can facilitate knowledge (at least in some cases). This is another way in which it would be wrong to think of the knowledge facilitating role of emotions (via marking/tagging information as important) as an evidential one.

*Highly Specific Appraisal Tendencies.* Many researchers have focused on how particular emotions produce different appraisal tendencies (even when two emotions are of the same valence).<sup>28</sup> (For discussion, see Lerner and Keltner, 2000) In particular, one such study (DeSteno et al., 2014) discovered that gratitude tends to reduce economic impatience, that is, it tends to "reduce excessive impatience for delayed rewards" and thereby combat one's tendency to "discount the value of delayed rewards relative to immediate ones". (The study also showed that sadness intensifies this unfortunate tendency). This effect on one's preferences will typically have a coordinate effect on one's normative judgments. In coordination with one's preference to wait, one will judge that it is better to wait. And when the feeling of gratitude is absent, one may be inclined towards a preference for immediate gain, and in making a decision, one might judge

---

<sup>28</sup> One should be skeptical of a valence generalization to the effect that any two kinds of emotions that are either positive or negative will produce the same kinds of effects or produce the same set of tendencies. See Lerner and Keltner (2000) and DeSteno et al., (2014) for discussion.

that it is ‘better to grab the reward that I can have right away rather than wait’ (a judgment that is reinforced by off-the-cuff sentiments like, ‘who knows what might happen between today and tomorrow’). If the study by DeSteno et al. (2014) is right, then gratitude can certainly play a safety enhancing role with regard to normative judgments about what is better to do.

Of course, it is possible that there will be cases where gratitude causes one to eschew immediate gain even when it is, in theory, more rational to opt for it. But recalling the phenomenon of *narrow relevant alternatives* (a crucial tool for the safety theorist), the existence of such cases will not preclude other normative beliefs caused by gratitude from being safe. And of course, it would be quite a stretch to suppose that gratitude is evidence for the decision-theoretic judgment (in those cases where gratitude facilitates safe normative beliefs about it being better to wait). This is especially clear when the role of emotions carries over to judgments that are not integrally related to those emotions in any way. This is known as the phenomenon of the *incidental carryover of emotions*. That is, emotions directed towards one event will have carryover effects on the shape of one’s thinking about unrelated matters. (For discussion, see Bodenhausen, 1993; Loewenstein & Lerner, 2003; and DeSteno et al., 2014). Consider an incidental carryover case where one’s gratitude towards someone facilitates a correct judgment that it is better to wait when making a decision that does not have anything to do with the original object of one’s gratitude. In such a case, it is especially clear that gratitude is nothing like evidence that it is better to wait.

*Emotions-as-Information.* When some epistemologists think about the epistemic role of emotions *vis à vis* judgments, they tend to gravitate to cases that fall into the category of what some psychologists call ‘feelings-as-information’ or ‘affect-as-information’. (For relevant

discussion, see Schwarz, 1990; 2011.)<sup>29</sup> The basic idea here is that one's judgments are directly guided by one's feelings towards a particular object, event, or even a theory.<sup>30</sup> In paradigm cases of this phenomenon, one is always conscious of the feelings and these feelings seem integrally (rather than 'incidentally') related to the subsequent judgment. As such, it is quite natural to speak of the judgment as being 'based' on the feelings. (We can contrast this idea with the gratitude case just mentioned—for despite the fact that gratitude plays an important causal role, it is less natural to speak of the normative judgment about it being better to wait as being 'based' on one's gratitude.) In many of the paradigm cases of 'affect-as-information', the output judgment is an evaluative one. (For example, consider a case where a judgment about whether things are going well is a direct response to one's emotional state.)

On such occasions, it is obvious that emotions have the potential to play a safety-inducing role. Suppose John feels optimistic and judges that his life will go well. John's optimism stably induces this belief that his life will go well and in fact, his life does go well at the world in question and in nearby worlds. Thus, John's belief is plausibly safe even by the standard of Super-Refined Safety. But in such a case, it is not obvious that John's optimism should be thought to play an evidential role. Rather, the most crucial aspect of John's optimism is in *stabilizing* his belief. Without optimism, John's judgment about whether his life will go well might be haphazard across close possible worlds and haphazdness precludes safety. Emotions might confer stability on hypotheses appraisals and as a result give judgments a shot at being safe. Anything that can confer stability can thus potentially enhance the capacity for safety.<sup>31</sup>

---

<sup>29</sup> The discussion here was inspired by a talk of Timothy Williamson's at Oxford (unpublished) where he sketched a way of thinking according to which the simplicity of a theory was not evidence for the truth of a theory but a stable preference for simple theories might nevertheless have a knowledge-conducive effect.

<sup>30</sup> For example, see Goldie (2004), Elgin (2008), and Helm (2009).

<sup>31</sup> But of course, conferring stability is not tantamount to probability raising.

Of course, I do not mean to preclude the possibility of ‘affect-as-information’ cases where it turns out the emotions do play an evidential role. (Indeed, the lesson of §1 and §2 is that there is no principled basis on which to preclude this.) But we should not think that every case in the broad ‘affect-as-information’ category (where emotions can help to explain safety by conferring stability) is also a case in which they play an evidential role. Consider a simple example.

Suppose John’s evidence is that which he knows and from the point of view of his ‘priors’ the prior probability that his life will go well is fifty-fifty. Suppose also that the conditional probability of John’s life going well, conditional on his feeling optimistic, is also only fifty-fifty. Finally, suppose that John is so constituted that feeling optimistic induces a belief that his life will go well. John feels optimistic and this produces the belief. Conditional on John’s feeling, John is only fifty-fifty that his life will go well. There is no probability raising of the proposition that his life will go well. And so, at least on the probability raising conception, John’s optimism is not evidence that his life will go well. Yet we can imagine that the belief that his life will go well is safe and this is partly explained by the fact that his feeling of optimism is stability-conferring.<sup>32</sup>

Admittedly, this sketch is no more than suggestive. But I hope to have done enough to show that it would be a mistake to focus exclusively on the evidential role of the emotions when inquiring about their positive epistemic value.

---

<sup>32</sup> Note the epistemic likelihood of a proposition depends on one’s total evidence and that includes not merely the fact that one is optimistic but also the fact that one’s life is going to go well (which becomes a part of one’s total evidence by being known).

### **3.4 Concluding Remarks**

Can emotions provide evidence for evaluative claims? I take myself to have provided a convincing case that they can and do. Yet the evidential role of emotions does not exhaust their epistemic worth. Another important role of emotions is as safety enhancers. If we ignore the many ways that emotions can figure non-evidentially in safe belief forming processes, then our understanding of their epistemic value will be radically deficient.

# Chapter Four

## Conditional Emotions

When conditionals embed factive emotional attitudes, they generate an apparent counterexample to modus ponens. In this essay, I present the puzzle and propose a solution—one that connects in interesting ways to the role of emotions in our suppositional reasoning and in our cognitive life more generally.<sup>1</sup>

### 4.1 The Challenge

Let us begin by setting the scene. Consider a simple conditional:

(1) If John is at the party then he is having a good time.

It is natural to assume that conditionals expressed in the English language obey modus ponens. (If a conditional is true, and the antecedent is true, then it follows that the consequent is true.) Thus, if it is true that John is at the party, and (1) is true, then it follows that he is having a good time. Meanwhile, if it is true that John is at the party, and it is false that he is having a good time, then the conditional, (1) is false.

We can generate an apparent challenge to modus ponens by inserting ‘know’ in the consequent. Suppose that I say,

(2) If John is at the pub then I know that John is having a beer.

---

<sup>1</sup> The puzzle I am concerned with is also presented in Drucker (2017) and Blumberg and Holguin (manuscript, 2018.), though neither opt for anything like my solution.

Let us suppose that given what I know about John (namely, that he always has a beer when he is in a pub), it is intuitive to think that the conditional (2) is true. Yet, it is important to notice that the most natural setting in which I utter a conditional like (2) is one where I do not know whether the antecedent is true, and so one where I do not know whether John is at the pub and do not know whether he is having a beer.

It is very easy to imagine a case where both the conditional (2) is intuitively true and the antecedent ('John is at the pub') is true. If we are right to assume that conditionals expressed in the English language obey modus ponens, then on the face of it, it seems to strictly follow that I must know that John is having a beer. This is a very strange and problematic result given that we have already stipulated that I do not know whether or not John is at the pub and that I do not know that John is having a beer. Given the details of the case, (2) seems to present an apparent counterexample to modus ponens.

The standard way to explain what is going on in the case of (2) is to say that the logical form of the conditional expressed by my utterance in (2) is not as it appears. In (2) the scope of the attitude verb 'know' appears to take narrow scope when positioned in the consequent of the conditional but at some level of deep structure, the logical form of my utterance is actually one where the attitude takes wide scope. Thus, the standard explanation posits that when I say (2), I really mean something more like (3).

(3) I know that *if John is at the pub, then John is having a beer.*

If this is the right way to think about the actual logical structure of my utterance in (2), then there is no counterexample to modus ponens after all. Notice that the consequent says

nothing about my knowledge anymore (it merely says that John having a beer). When a conditional like (2) is true, the apparent challenge posed by a narrow-scope reading disappears under the wide-scope reading. In sum, we use the wide-scope maneuver to show how embedding knowledge in the consequent of a conditional like (2) poses no actual threat to modus ponens after all.

### ***The Phenomenon***

Factive emotional states are a special class of attitudes that pose a similar challenge to modus ponens and in these particular kinds of cases, the solution is not straightforward. Suppose that we are in a setting where the following conditional (4) is intuitively true and I lack knowledge of the truth value of the antecedent:

(4) If John is at the pub, I'm angry.

Preserving the assumption that conditionals obey modus ponens, it follows that—even though I do not know whether John is at the pub, I'm angry. (4) is an example of a very strange phenomenon that occurs when we embed factive emotional attitudes into the consequent of conditionals. It is very natural to think that I am not angry in the situation where I utter (4) given that I do not know whether John is at the pub. (Indeed, assuming that the consequent is shorthand for, 'I'm angry that John is at the pub,' the claim that the consequent is true, and means what it seems to mean, violates the very plausible claim that being angry that P entails knowing that P.)<sup>23</sup> Thus, (4) seems like another *prima facie* counterexample to modus ponens.

---

<sup>2</sup> This consideration applies directly when we focus on sentences like, 'If John is at the pub, I'm angry that he is.'

<sup>3</sup> Drucker (2017) argues that I am angry that John is at the pub in these settings. His idea is to learn to live with a radical externalism about attitudes like anger – one can end up being angry at someone even if one

The standard solution in the case of ‘know’ is to make the wide-scope move (adjust the narrow-scope of the attitude verb by moving it from its position in the consequent to the front of the antecedent.)<sup>4</sup> Recall that the wide-scope move is meant to give us a more accurate reading of the true meaning of the conditional by appealing to deep logical structure.

When we adjust the scope of my attitude in (4), we get the following result:

(5) I'm angry that if John is at the pub, then \_\_\_\_.

As (5) shows, the wide-scope move immediately presents us with a new problem: there is nothing left in the consequent. We can repair this initial problem by filling in the conditional—so let's suppose that what I meant by (4) is (6). Adjusting the scope gives us (7) as result.

(6) If John is at the pub, then I'm angry that John is at the pub.

(7) I'm angry *that if John is at the pub, then John is at the pub*.

The result generated by the making the wide-scope move gives us (7) which sounds like an absolutely terrible thing to say.<sup>5</sup> The first thing to notice is that anything of the form 'if P, then P' is a tautology.<sup>6</sup> Moreover, it does not seem like tautologies are the sort of thing for me to be

---

has no credential states about that particular individual. Drucker never addresses the fact that this idea clashes with a knowledge constraint on factive emotional constructions. Blumberg and Holguin (manuscript, 2018) agree with me that ‘I am angry that P’ entails I know P. They also agree that on the face-value reading, ‘I am angry that John is at the pub’ is false in the contexts I am concerned with. Their solution is rather different from what I propose in section two as it appeals to a non-standard semantics for these constructions. I discuss their solution in the appendix.

<sup>4</sup> A *locus classicus* for discussion of scope is Russell's, ‘On Denoting’ (1905).

<sup>5</sup> The inadequacy of wide-scoping is also emphasized in Drucker (2017) and Blumberg and Holguin (2018 manuscript).

<sup>6</sup> Cf. Drucker (2017), who also points this out.

angry about. This is even more vivid in the case of surprise. Consider the following minimal pair:

(8) If John is at the pub, then I'm surprised that John is at the pub.

(9) I'm surprised that *if John is at the pub, then John is at the pub*.

(8) seems to be a perfectly fine thing to say (though it is an instance of the phenomenon) but like (7), (9) also sounds terrible because it is completely obvious that tautologies are not surprising. In sum, the wide-scope move that we standardly use in the case of 'know' is not available for all other conditionals that involve factive attitudes (as is clear from the case of embedding factive emotional attitudes in conditionals like 4, 6, and 8). But if the standard explanation cannot make sense of these particular kinds of cases, we might ask: Are factive emotions (when embedded in conditionals like 4, 6, and 8) an actual counterexample to modus ponens? No. Probably not. Recalcitrant puzzles are rarely a sound basis on which to revise established principles of logic. But we are now left with the task of making sense of the phenomenon in a way that preserves classical logic.

## 4.2 The Cognitive Explanation

My explanation of the phenomenon will deploy some of the tools given by Timothy Williamson's work on conditionals and supposition (2016 and forthcoming). Let's begin by highlighting some key ideas from his discussion.

The way that we ordinarily evaluate conditionals is very similar to the way that we update on new information.<sup>7</sup> Consider the conditional, 'If the train is delayed by fifteen minutes, I am running late to work.'

---

<sup>7</sup> Cf. Ramsey, 1929 and Williamson, 2016. See Chalmers and Hajek (2007) for some concerns about Ramsey's idea (concerns that actually connect with this discussion).

There is a possible scenario where I learn that the train will be delayed by fifteen minutes and when I update on this new information, I conclude that I am running late to work. Now consider a suppositional case. Suppose I am waiting for the train. I notice that the train is running a few minutes late and I do not know how much longer the train will be delayed. As I am waiting, I begin to run through the various imaginary scenarios where the train still has not arrived in two minutes, and in three minutes, and in fifteen minutes. The basic thought is that when I suppose that the train is delayed by fifteen minutes, I arrive at the same conclusion that I would arrive at in the possible scenario where I learn that the train is delayed by fifteen minutes, and I conclude (under the supposition) that *I am running late to work*.

Under the supposition that the antecedent is true, I simulate the way that I would reason in the actual case and this offline-reasoning process leads me to offline-believe the consequent which, in turn, generates an online-belief (a genuine belief) in the conditional, ‘If the train is delayed by fifteen minutes, I am running late to work.’ In general, this is one way we use supposition to evaluate conditionals.

**Observation 1:** *When reasoning under supposition, we typically reason using the following heuristic: ‘If I offline-believe that  $P$  under supposition  $S$ , I believe the conditional: If  $S$  then  $P$ .’*

Under supposition, we reason in a way that closely approximates the way that we actually reason when we update on new information. When one simulates their reasoning under supposition, their reasoning is ‘offline’ and culminates in ‘offline-states’. Thus, when I reason under the supposition that the train is delayed by fifteen minutes, I come to offline-believe

exactly what I would believe in the situation where I learn that the antecedent obtains. This offline-belief is not ‘online’ because its input is suppositional and not real-world learning.<sup>8</sup>

Although offline-reasoning under supposition is inferentially similar to the way we reason online (e.g. when Tim believes/offline-believes P&Q, Tim believes/offline-believes P and Tim believes/offline-believes Q) offline-reasoning is likely to be disconnected from intentional action in some important respects. If I learned that the pig pen has been left open, I will conclude that some of my pigs are loose and this, in turn, will prompt various actions: Rushing to the garden, calling for John to capture the loose pigs, etc. Consider the case where I suppose that the pig pen has been left open. Now, I will offline-conclude (i.e. form the offline-belief) that some of my pigs are loose. But this offline-belief will not (generally) culminate in analogous actions. While there is an analogy in inferential role, offline-beliefs do not engage with my desires and intentions in a way characteristic of genuine beliefs. (Of course, it is wrong to say that offline-beliefs are totally disconnected from action. For example, offline-beliefs may play a role in getting me to online-believe some conditionals which, in turn, might yield various actions.)<sup>9</sup>

Of course, the parallel between what we offline-believe under a supposition and what we would believe if we learned the antecedent of the conditional is not exact. For example, if I suppose that someone is stealing from me without my ever learning about it, it would not make sense to say that whatever I offline-believe under that supposition is exactly what I would believe if I learned that *someone is stealing from me without my ever learning about it*. Why? Because that is not something I could learn. Consider another example. If I suppose that a bear is running

---

<sup>8</sup> ‘Very roughly, the online and offline processes take the same input... and deliver the same output...by the same means. One process is online and the other is offline in virtue of the different sources of the input.’ Williamson (2016, p. 118)

<sup>9</sup> Of course, if you are concerned enough about loose pigs, then having believed the conditional, you might check the pig pen anyway.

towards me right now, it would be natural for me to develop the supposition by offline-believing that I don't know it, yielding the intuitively correct conditional: 'If a bear is running towards me right now, I don't know it.' But if I learned that a bear was running towards me, I would not think, 'I don't know it.' In that respect, what I offline-think under the supposition does not (in this case) match what I would think if I learned the antecedent.

Returning to the phenomenon at issue, how can we explain what is going on in cases where conditionals embed factive emotional attitudes? We can answer this question by first observing that when we suppose things, we do not merely form offline-beliefs but we also form offline-emotions. Just as we simulate what we would believe (offline) under a supposition, we also simulate the way we would feel (offline) in a world under that supposition. (For example, under the supposition that I will be burned at the stake tomorrow night, I'm offline-terrified.) This observation will play a key role in the cognitive story of what is going on in cases of conditionals that involve emotions.

**Observation Two:** *When reasoning under supposition P, we typically have various offline-emotions such that for various emotional V's, we offline-V that P.*

The pattern described in the case of belief/offline-beliefs carries over to emotions/offline-emotions. Just as there are deep analogies between what we believe upon learning P and what we offline-believe upon supposing P, there is a deep connection between what emotions we would have upon learning P and what emotions we would offline have upon supposing P.

It is also worth mentioning that offline-emotions are disconnected from action in a way that is similar to how offline-beliefs are disconnected from action. Under the supposition that I will be burned at the stake tomorrow, I might have all kinds of involuntary physiological

reactions (involuntary reactions that are less intense versions of the ones I would likely have if I *actually* learned that I will be burned at the stake tomorrow).<sup>10</sup> Perhaps my face turns pale, I cringe, I feel my stomach turn, etc. And the intentional content of my offline-terror is similar to the content of the terror in the case where I learn that I will be burned at the stake tomorrow. But in the actual situation where I learn this, my terror will prompt me to beg for mercy, and make plans of escape. Under the supposition that I will be burned at the stake tomorrow, my offline-terror does not cause me to act in those ways.<sup>11</sup>

The cognitive explanation of the phenomenon requires a third and crucial observation. Ordinarily, when we learn that P and subsequently V that P (e.g. feel angry/surprise/happy/etc. that P) we also believe that we V that P. When I see that Tim is dressed like a cowboy, I am surprised that Tim is dressed like a cowboy, and I believe that I am surprised that Tim is dressed like a cowboy. Similarly, under the supposition that Tim is dressed like a cowboy, I am offline-surprised that Tim is dressed like a cowboy, and I offline-believe that I am surprised that Tim is dressed like a cowboy. In short, we can observe that there is a transition from having an offline-emotion under a supposition to having an offline-belief about having that particular emotion.

**Observation 3:** *The offline-emotion to offline-belief transition: When we offline-V that P (for emotional verbs) under a supposition that P, we tend to also to have an offline-belief that we V that P.*

---

<sup>10</sup> ‘There is a record showing the vasomotion of anger, a marked rise, through purely imaginary anger. I do not have to get my subjects really angry; merely to have them imagine that they are angry, recall certain company etc. is enough to have a great rise in the systolic and diastolic pressures.’ Dearborn (1918, p. 20.)

<sup>11</sup> Note in passing that the phenomenon of offline-emotion has some capacity to shed light on the much-discussed topic of fearing fictions. (See Walton, 1978.) The kind of offline-sadness that I have when I suppose that a chef has cooked my pig into a schnitzel is similar to the mental state I undergo when I watch a movie in which someone’s pet pig gets schnitzeled. In both cases, the temptation to call the mental state ‘genuine sadness’ is encouraged first, by a tendency to think of sadness as merely a quale, and second, by a tendency to think of the involuntary physiological reactions associated with sadness as sufficient for being sad.

It is important to be very clear that the idea is *not* that under the supposition, I offline-believe that I am offline-angry/surprised/etc. If we are trying to provide the best cognitive explanation of what is going on, we must remember that most people do not have beliefs (online or offline) about ‘offline-emotions’ (this is really just a term of art.) The key idea is that when we suppose P, and then offline-V that P, we also tend to have an (offline) belief that we V that P. In other words, an offline-emotion under a supposition tends to trigger an offline-belief about oneself as having an online-emotion.

We can now make sense of the phenomenon at issue in light of the discussion above. Suppose John promised me that he would work on the farm today instead of going to the pub. Around noon, Jason (our farm hand) calls me and asks me about John’s whereabouts.

*“Well, John promised that he would be working on the farm today. Maybe John is working in the stable. Go check, and call me back if you don’t find John there.”*

I hang up the phone and wonder where John is. In my mind, I play out two scenarios. In the first, I suppose that Jason finds John working in the stable, and under that supposition, the world that develops in my mind is one where I’m offline-relieved that Jason found John working in the stable, and I offline-believe that I’m relieved that Jason found John working in the stable. I also imagine a scenario where it turns out that John is at the pub. In the scenario where I suppose that John is at the pub, I’m offline-angry that John is at the pub (observation two), and I offline-believe that I’m angry that John is at the pub (observation three.)

The cognitive explanation of the phenomenon generated by the conditional ‘if John is at the pub, I’m angry that John is at the pub,’ can be made vivid by considering how the three

observations figure in the story above. (An analogous story can be told about the conditional, ‘If John is in the stable, I’m relieved.’)

When I begin by supposing that John is at the pub, observation two tells us that: (i) I’m offline-angry that John is at the pub. Observation three tells us that: (ii) I offline-believe that *I’m angry that John is at the pub*. And the heuristic we use to evaluate conditionals that is described in observation one generates my online-belief in the following conditional:

(6) If John is at the pub, I’m angry that John is at the pub.

We know that conditionals obey modus ponens and yet these cognitive habits (as detailed by the three observations above) work together to actually produce a false judgment about the conditional (which thereby generates the phenomenon: an apparent counterexample to modus ponens). In summary, the idea is that the way we ordinarily evaluate conditionals leads us to think that conditionals like (6) are true. But (6) is not a conditional that is true. We make mistakes about conditionals when they involve factive emotional states because the heuristic we ordinarily use to evaluate conditionals leads us to mess up in these particular kinds of cases. There is no problem for modus ponens after all and the phenomenon disappears. In short, modus ponens isn’t wrong—we are. Our cognitive habits lead us to make mistakes when evaluating these particular kinds of conditionals—and in such cases, we are simply not reasoning carefully.

The idea that this heuristic (described by observation one) often leads us into error is not an objection to my proposal. After all, Williamson (forthcoming) has described various cases (that are quite different from this one) where this heuristic gets us into trouble.<sup>12</sup>

---

<sup>12</sup> In Williamson (forthcoming), he suggests that it gets us into trouble with counterpossibles and also in our estimation of the probabilities of conditionals.

### 4.3 Some Natural Language Data

Let us consider three pertinent pieces of natural language data that the cognitive explanation does well to explain.

#### The First-Person/Third-Person Data

One might have already noticed that the challenge posed by the phenomenon (generated by embedding factive emotional attitudes in conditionals) has thus far been presented using examples where the consequent is in the first-person. Does the phenomenon occur with consequents in the third-person? Consider the following minimal pair.

(4) If John is at the pub, I'm angry.

(10) If John is at the pub, John's mother is angry.

Suppose that I know I do not know whether John is at the pub and I know that his mother does not know whether John is at the pub. Recall that it is intuitive to think that (4) sounds fine even if I am in a setting where (i) John is at the pub and (ii) I do not know whether John is at the pub. We know that (4) is an instance of the phenomenon and we know that it sounds like the sort of thing that we would say in the relevant contexts but (10) has importantly different felicity conditions.

In the third-person, (10) is felicitous if John's mother is in a context where I expect that if it is true that John is at the pub, then his mother will know it. In that context, (10) does not present a challenge to modus ponens. But in the context where I do not expect that John's mother knows whether John is at the pub, (10) sounds like a terrible thing for me to say. In short, (10) is not intuitively true in a setting where (i) John is at the pub, and (ii) John's mother does not know whether John is at the pub.

*The Cognitive Diagnosis.* The cognitive explanation can make sense of the contrast brought out by the first-person/third-person data. Observation three details the transitional mechanism whereby one goes from having an offline-emotion under a supposition to having an offline-belief about having an online emotion. This transitional mechanism is not available in the third-person case.

There is no such transition from my having an offline-emotion to having an offline-belief that John's mother or anyone else (but me) has an emotion. Of course, in a setting where I know that John's mother knows whether John is at the pub (and I know that is the sort of thing that would make her angry if it were true), my offline-supposition will induce an offline-belief that John's mother is angry but that would not be triggered by my own offline-emotional state.<sup>13</sup>

### **The 'Right-Now' Data**

Consider the following minimal pairs.

- (6) If John is at the pub, I'm angry that John is at the pub.
- (11) If John is at the pub, I'm angry right now that John is at the pub.
- (6) If John is at the pub, I'm angry that John is at the pub.
- (12) If John is at the pub, right now I'm angry that John is at the pub.

---

<sup>13</sup> There is a slightly different phenomenon where the first-person/third-person contrast is less important. In a situation where I do not know whether anyone has contributed to my charity, it wouldn't be unnatural to say, 'If anyone has contributed to my charity, I'm very grateful.' (I owe this example to Tim Williamson.) But in this kind of case, it is not so bad either for someone to say of me, 'If anyone has contributed to her charity, she is very grateful.' Similarly, 'If anyone is doing charity-work in that war zone, I admire them' and 'If anyone is doing charity-work in that war zone, she admires you.' As Williamson pointed out in conversation, this phenomenon is not confined to conditionals: 'I admire everyone/anyone doing charity-work in that war zone' is also felicitous. I would like to tentatively suggest that this turns on a liberal use of 'admire' and 'is grateful.' After all, in this kind of setting, we can also get away with, 'She has considerable admiration for you' when addressing a group of charity workers in the war zone. (By contrast, 'She is angry that John is at the pub' is completely unacceptable in the original example.) Examples similar to the gratitude and admiration examples figure alongside anger and regret examples in Drucker (2017) – Drucker does not distinguish them properly from the phenomenon that I am primarily concerned with.

Notice that (11) and (12) sound terrible but (6) does not. Why is this?

*The Cognitive Diagnosis.* The cognitive explanation that I have suggested explains why we assert (6) and can explain why we do not similarly assert (11) and (12) via the same mechanisms: The reasons why we make mistakes in conditionals like (6) but not in conditionals like (11) and (12) is because when we evaluate (11) and (12) ‘right now’ draws attention to the fact that I am not (in reality) angry that John is at the pub (because I do not know whether John is at the pub.)<sup>14</sup> Moreover, given that we know (i) conditionals obey modus ponens and (ii) for all we know, the antecedent is true—when our attention is drawn to the fact that the consequent is false, we are blocked from asserting the conditional. When we get sloppy and forget that the consequent is false, we assert conditionals that generate the phenomenon (e.g. ‘If the chimney cleaner saw me naked, I’m horrified that he did.’)

### **Finite Conjunction**

The cognitive explanation can also make sense of a similar kind of infelicity that is generated by finite conjunction cases which involve conditionals that embed factive attitudes. As a principle of reasoning, finite conjunction tells us that:

P1: If P then Q.  
 P2: If P then R.  
 C: If P then (Q&R).

But consider the following case:

P1: If the chimney cleaner saw me naked, I don’t know that he did.  
 P2: If the chimney cleaner saw me naked, I’m horrified that he did.  
 C1: If the chimney cleaner saw me naked, I don’t know that he did and I’m horrified that he did.

---

<sup>14</sup> Consider the placement of ‘right now’ in the conditional, ‘If John is at the pub, I’m angry that John is at the pub right now.’ In this case, the conditional doesn’t sound too bad and this is because ‘right now’ is not drawing attention to the fact that I am not angry at the time of my asserting the conditional. In this conditional, ‘right now’ specifies the time of John’s being at the pub, not the time at which I am angry.

C2: If the chimney cleaner saw me naked, I'm horrified that he did and I don't know that he did.

Individually, the conditionals in P1 and P2 sound felicitous but C1 and C2 sound radically infelicitous. Yet finite conjunction tells us that each C1 and C2 follow from P1 and P2. Presumably, we want to preserve finite conjunction so this case calls for some explanation.

This infelicity is explained by the fact that 'I don't know that he did' is playing the same role in C1 and C2 as 'right now' played in (11) and (12). It is vivid to us that the consequent is false. It is obvious that you cannot be horrified that P while not knowing that P. Given that the consequent is false and modus ponens is true, we are blocked from asserting C1 and C2. The cognitive habits that allow us to be sloppy and assert P2 do not lead us to assert C1 and C2.

#### **4.4 Concluding Remarks**

To end, we might ask whether the cognitive explanation can also make sense of conditionals that are thought to require the standard wide-scope maneuver (as discussed in section 1). The standard explanation says that the true meaning of such conditionals is represented by adjusting the scope of the attitude in the conditional. Perhaps the cognitive explanation can do better because it can do more explanatory work. Perhaps what is actually going on is that our reasoning patterns are uniformly leading us to make mistakes. After all, the wide-scope move is not available for all conditionals with embedded attitudes (as demonstrated by the fact that wide-scope move was not available for conditionals that embed factive emotional attitudes.)

The cognitive explanation offers a different account of what is going on in cases where I assert (2) If John is at the pub, I know that John is having a beer. The cognitive diagnosis of (2) says that under the supposition that John is at the pub, I offline-know he is having a beer, and I

also offline-believe that I know that he is having a beer. In other words, we can extend observation three to other attitudes such as believing and knowing. Although I will not explore or defend this idea further here, this approach should be taken seriously given its power to explain the phenomenon generated by conditionals that embed factive emotional attitudes—a phenomenon that was not resolved by the standard wide-scope maneuver.

***Appendix. Blumberg and Holguin's Semantic Approach***

I have proposed a cognitive explanation according to which the relevant conditionals come out false but are nevertheless asserted in special situations where our cognitive habits lead us into trouble. Blumberg and Holguin (manuscript, 2018) propose an alternative account which develops a very sophisticated semantics of these conditionals in order to avoid an error theory.

Blumberg and Holguin take it that conditionals that embed ‘know’ in the consequent can come out true in the relevant cases without relying on the standard wide-scope move.<sup>15</sup> They work with a fairly standard picture where, on a face-value reading, ‘I know P’ is true just so long as there are no epistemically possible worlds where  $\neg P$ . They propose that on a special Q-restricted reading, ‘I know P’ is true just so long as there are no epistemically possible Q-worlds where  $\neg P$ . Their basic idea is that when, ‘I know P’ occurs in the consequent of a conditional, there is an available reading according to which, ‘I know P’ is restricted by the antecedent. For example, consider the conditional:

(13) If John is not at work, I know that John is at the pub.

---

<sup>15</sup> Note that my brief summary of Blumberg and Holguin's account uses less formalism because it is not necessary for the points I wish to make here.

If I utter (13) in a setting where I do not know whether John is at work then it is not true that at all epistemically possible worlds he is at the pub (e.g. John might be at work). But it might be that at all epistemically possible worlds where John is not at work, he is at the pub. Thus, the (13) can come out as true without making the wide-scope move. Building on this, the authors suggest a restricted account of emotional verbs. On a Q-restricted use, ‘I regret P’ is interpreted along the same lines as one standard analysis of ‘regret’ except one uses a Q-restricted version of ‘know’ in the analysis.<sup>16</sup>

While I will not go into the details of their specific proposals about the semantics for various factive emotional attitudes and restricted variants, I think it is worth raising a few general concerns that I have about their program.

*i.* Blumberg and Holguin’s semantic approach posits a fancy restricted semantics for various emotional verb/adjective + complementizer phrases (e.g. ‘regrets that P’ and ‘is surprised that P.’) But it is important to be aware of the fact that there is an extensive range of complementizer-phrase constructions that display the phenomenon of interest. Consider, for example:

(14) If John is at the pub, I’m surprised and have a nagging feeling that he will be home late.

---

<sup>16</sup> Blumberg and Holguin (manuscript, 2018. p. 12-13) are inspired by a gloss on ‘regret’ offered by Irene Heim: ‘S regrets P’ is true at W *iff*: (i) S knows P at w and (ii) for most worlds w’ compatible with what S knows at w: S prefers (at w) the worlds most similar to w’ in which  $\neg p$ , to w’. (One gets a Q-restricted gloss on regret by Q-restricting each occurrence of ‘know’ in the definition. See MS. p.13) I note in passing that Heim’s analysis is implausible. Suppose I make a bet and don’t know the outcome. Suppose it is likely that I will make a small loss but there is a chance of a huge gain. Most worlds where I make the bet, I lose and for each of those worlds, the closest worlds where I don’t bet is preferable. But that doesn’t mean I regret the bet – it is easy to flesh out the details where I am actually very satisfied with the bet (by making the size of the gain compensate for its relatively low probability). Thanks to John Hawthorne for extensive discussion here. (That said, Holguin and Blumberg are using Heim’s analysis for illustrative purposes only.)

(15) If John is eating dinner alone, I feel a little sad that he didn't invite me to eat with him.

It remains to be seen whether there is a plausible restricted reading for, 'having a nagging feeling that P' or 'feeling a little sad that P' and so on. (I am skeptical.)

*ii.* While Blumberg and Holguin's semantic approach offers an account of the emotional verb+complementizer phrase constructions, the phenomenon is not confined to such constructions. In fact, we can generate the phenomenon by having a conditional where the emotional attitude takes a noun phrase complement. For example, consider, 'If Jason left the pig pen open, I'm annoyed by him,' and, 'If Jason is the pig-killer, I feel nothing but intense hatred for him' (both are felicitous utterances in suitable contexts.) The cognitive explanation extends to these conditionals without any problem—but it remains to be seen how the semantic approach will fare.

*iii.* The semantic approach has no useful account of why the phenomenon arises in the first-person case and not in the third-person case. Although Blumberg and Holguin notice the first person/third-person contrast (as well as a contrast between the present tense and other tensed cases) they do not offer any explanation that follows from any semantic or conversational principles. Instead, the authors introduce a special rule according to which one is only allowed to use 'Q-restricted readings on attitudes when the clause is present tense and in the first person'<sup>17</sup> This strikes me as a rather unsatisfying move.

*iv.* The semantic approach does not offer a useful account of the 'right now' data. In those cases, the conditionals are infelicitous despite the fact that the relevant consequents are in the first-person and present tense. If the restricted reading was available for first-person present tense

---

<sup>17</sup> This is their 'First-Present' restriction, see Blumberg, & Holguin (manuscript, 2018, p. 23).

consequents, it would not sound infelicitous. I do not see any good explanation of this data forthcoming from the semantic account.<sup>18</sup>

v. Lastly, it is worth considering conditionals with consequents that display the phenomenon where there is considerable pressure to read the emotional verb at face-value.

Consider the following exchange:

*Adams: 'I hate John!'*

*Bill: 'If he's where I think he is, I do too!'*

Obviously the first use of 'hate' is not restricted. Given 'do too' picks up on that, there is reason to think it should be interpreted in the same way (i.e. not restricted.)

I should like to emphasize that I do not want to discount the semantic approach entirely but I do think the challenges outlined above give us reason to think my cognitive explanation may have more promise.

---

<sup>18</sup> Note that the semantic approach does better with the finite conjunction data so long as it can make plausible that if one Q-restricts one attitude in the consequent one should Q-restrict others. 'If John is in the pub, I don't know it' is false on the Q-restricted reading, true on the non-restricted reading. Meanwhile, 'If John is at the pub, I am angry that he is' will presumably come out false on the non-restricted reading but true on the restricted reading. Some suitable uniformity requirement will block the availability of a good reading for 'If John is at the pub, I don't know it and am angry that he is.'

## Chapter Five

## Doxastic Cognitivism

Cognitivism about the emotions is typically associated with the view that emotions are evaluative judgments.<sup>1</sup> This particular form of cognitivism is notably represented by Robert Solomon and Martha Nussbaum. Solomon (1977, p. 125) writes, ‘An emotion is a judgment (or a set of judgments) ... about our Selves and our place in our world.’ Nussbaum (2004) develops similar ideas: ‘My sadness, my sorrow, and my grief are judgments of various severity to the effect that I have suffered a loss.’ And she notably argues for the general thesis that emotions are, ‘forms of evaluative judgment ...’ (2001, p. 22)

Judgment-based cognitivism (as represented by Solomon and Nussbaum above) has inspired a spirited discourse among cognitivists and non-cognitivists alike. A number of theorists have taken issue with the idea that emotions are judgments while maintaining that emotions still somehow represent evaluative states of affairs.<sup>2</sup> (This has led to a proliferation of alternatives to the judgment-based cognitivist proposal.)<sup>3</sup> In this discussion, I present and develop an under-explored alternative to traditional cognitivist approaches, an alternative that I call ‘Doxastic Cognitivism’.

Cognitivists are often accused of over-intellectualizing the emotions. I hope to show that this criticism only arises when the cognitivist is not careful to distinguish judgment from belief, and when theorists posit an over-intellectualized notion of belief. The view that emotions involve evaluative beliefs need not be over-intellectualized. Doxastic Cognitivism (as I develop it) is a

---

<sup>1</sup> ‘Cognitivism’ is a term of art. It will not serve our purposes to have a precise definition here.

<sup>2</sup> Another possible point of departure concerns whether the contents are evaluative. The Jamesian thesis that emotions are perceptions of bodily states offers an interesting alternative to the evaluative content approach. See James (1884). I will not engage with this approach in detail here.

<sup>3</sup> For a survey of cognitivist views, see Deigh (1994) and Scarantino (2010).

novel and robust cognitivist theory, and I hope to show how a strong positive case can be made for the view.

## 5.1 Preliminary Distinctions

Let us begin by noting a few important structural details that will shape our discussion.

Cognitivism is often advanced with little regard to the distinction between belief and judgment. Yet, it is very important to be clear about whether the cognitivist's central idea is that *emotions involve evaluative beliefs* or whether *emotions involve evaluative judgments*. The distinction between belief and judgment is fairly standard. A belief is a mental state and a judgment is an act that the agent performs. (It may be useful here to think of the act of *judging* as something like an internal act of assertion.)<sup>4</sup> We also find this distinction in ordinary language. While it is very natural to talk about judgments in the progressive, it is never natural to talk about beliefs in the progressive. For example, consider how unusual it sounds to say that *John is believing that Frank is guilty*. It is more natural to say that John believes that Frank is guilty.

Beliefs and judgments differ in another important respect: One might have millions of beliefs at any given time. For example, I believe that my name is Christina, I believe that Socrates was a philosopher, I believe that my coffee is getting cold, etc., but there is something peculiar about supposing that agents make millions of judgments at any given time, and this seems to be because judgments are more episodic than belief.

---

<sup>4</sup> The Stoics drew a distinction between an impression that something is the case (a *phantasia*) with a judgment that something is the case (*krisis*). In their thinking, the mental act of judging that P is an exercise of a sufficiently developed reasoning faculty (*logos*). The exercise of this faculty allows us to make judgments about the way things really are in light of our impressions (which tell us of the way that things seem). For a more extensive discussion, see Sorabji (2000).

Some cognitivists are not careful about the distinction between judgment and belief. For example, Nussbaum vacillates between the two ideologies. Although she tends to gloss emotions as judgments, in a paper entitled '*Emotions as Judgments of Value and Importance*' Nussbaum writes:

. . . emotions embody not simply ways of seeing an object, but beliefs – often very complex – about the object. .... In order to have anger I must have an even more complex set of beliefs: that there has been some damage to me. (p. 188)

It is also important to be sensitive to the difference between feeling a particular emotion (feeling-V, where V stands for some emotional verb) versus being in a particular emotional state (being-V). While a feeling of anger is often thought of as a state that one undergoes when one has a certain kind of conscious bodily disturbance, not all emotions are so experienced. Like beliefs, one might have many emotions about things without the associated feelings (e.g. being angry without feeling angry, being afraid without feeling afraid, etc.)<sup>5</sup> Consider a simple case where someone goes through life hating their enemy. It is not as though retaining hatred for one's enemy requires that one experiences a constant stream of those feelings associated with hatred. I suspect that those who identify emotions with judgments are focusing on events of feeling emotion rather than the states of having emotion. Thus, for example, the view that jealousy embodies a judgment that *someone else has something that oneself lacks* is much more natural if one is thinking of a case where an agent merely has some episode of feeling jealous rather than states of being jealous. For example, suppose Jones is jealous of his older brother's success.

---

<sup>5</sup> Note that I am not talking about the use of 'angry' to pick out a character trait, as in one central use of 'he is an angry person.'

Jones' state of jealousy might fester in the back of his mind for years even though his feelings of jealousy are merely episodic (say, whenever Jones' thinks about his brother's luxurious lifestyle).

When theorizing about the emotions, it is very important to be sensitive to the differences between belief and judgment and feeling-V and being-V. Insofar as we are theorizing about the states of being angry, being happy, being jealous, etc. as opposed to their felt manifestations, the above discussion suggests that a belief-oriented cognitivism is more plausible than a merely episodic judgment-based cognitivism.

On a final preliminary note, it is worth emphasizing that there is an important difference between the evaluative content of some emotion (the content of the purported belief or judgment that the emotion involves) and the content given by the that-clause in ascriptions of the form, 'S V's that P' (where S stands for the agent, V stands for the emotional state, and P for some proposition). For factive emotion ascriptions, we can gather from the truth of the ascription that (i) P is true, (ii) S knows that P, and that (iii) one of S's reasons for Ving is P.<sup>6</sup> If John regrets that P then it follows that John knows that P and that P is one of his reasons for regret. Any theorist must be careful to notice that the evaluative content given by the emotion itself does not require truth, knowledge, or any other relation to reasons. (If the regret is an evaluative belief of some sort, that belief is not naturally thought of as a reason for the regret.) It is also important to notice that the proposition given by the that-clause is not necessarily evaluative. 'The sun is out' does not have any evaluative content and yet one might be very happy that the sun is out. As a rule of thumb, one should not attempt to locate the evaluative content of any particular emotional state in the ordinary ascription of the emotional state (whether it is 'S V's that P' or in a case where an agent's emotional state is described by a claim of the form, 'S V's noun-phrase.')

---

<sup>6</sup> For further discussion, see chapter two.

In an effort to keep things under control, our discussion will focus on cases of being-V and the views under consideration will fall under the heading of what I shall call, ‘Doxastic Cognitivism’. For now, we can gloss Doxastic Cognitivism as the thesis that emotions involve an evaluative belief that P (not a judgment that P). For example, fear involves the belief that something is dangerous. (We shall be more precise in due course.)

## 5.2 Anti-Intellectualism

The Doxastic Cognitivist is an anti-intellectualist about propositional attitudes (like knowledge and belief) and subscribes to a fairly standard notion of propositional attitudes. According to the most standard picture of propositional attitudes, a sentence of the form, ‘S Vs that P’ is a claim of the individual (picked out by S) that S stands in a relation (picked out by V) to the proposition, P. Propositions are mind-independent, atemporal objects that are either true or false regardless of whether anyone ever stands in any sort of relation to them. Following Frege (1892), we can think of propositions as objects existing in a metaphysical third realm though the metaphysical details are not important here. What is important is that when S V’s that P, the proposition is not an item located in S’s head. Accordingly, one and the same proposition can be the object of numerous people’s attitudes (e.g. many people can believe the same proposition at the same time).

Despite this very standard way of thinking about propositional attitudes, there are still competing pictures about how cognitively sophisticated one needs to be in order to stand in attitudinal relations to propositions (e.g. for S to know that P or for S to desire that P).<sup>7</sup> In this connection, it is very important to distinguish an intellectualist approach from an anti-intellectualist approach. The intellectualist picture-thinker roughly has the idea that when S V’s

---

<sup>7</sup> Even rocks can stand in relations to propositions (e.g. the relation of numerical distinctness). Indeed, everything can stand in a relation to a proposition but here, I am only interested in attitudinal relations.

that P, there is something like a sentence, or a string of words, or ‘concepts’ in their head. The intellectualist is thinking that when S believes that P, a sentence-like representation of P gets filed away in a belief-box.

This is a crude but surprisingly popular way of thinking about propositional attitudes (encouraged by Fodor).<sup>8</sup> Even when theorists do not literally think of propositional attitudes as being based on sentences in the head, it is commonplace to talk of propositional attitudes as being mediated by, ‘mental representations’ that are built from ‘concepts.’ (And the structure of such ‘representations’ and ‘concepts’ bear a striking resemblance to linguistic structures like sentences and noun phrases.)

The anti-intellectualist will concede that language can provide ways to stand in the propositional attitude relation (i.e. accepting sentences, or using sentence-like representations of propositions.) But the anti-intellectualist does not think that one needs anything like language to stand in an attitudinal relation to a proposition. It is wrong to think that propositional attitudes (i.e. knowledge, belief, desire) or even cognition in general, requires the use of something like linguistic-vehicles or concepts. Why? Because there are other ways to stand in an attitudinal relation to a proposition. Of course, when we say something of the form ‘S believes that P’, we attribute a belief to S by using a sentence in the that-clause to characterize the content of the belief, but this hardly implies that S needs to have a sentence-like vehicle in mind (that corresponds to the sentence in the that-clause) for our belief-attribution to be true. Robert Stalnaker writes:

...while theoretical and experimental developments in cognitive psychology may someday convince me that I store beliefs in a form that is structurally similar to the form in which they are expressed and described in English, I don’t think my ordinary belief attributions commit me to thinking that they will. (1990, p. 230)

---

<sup>8</sup> See Fodor (1978, p. 519)

Williamson and Stanley's account of knowledge-how (2001) is one paradigm of the anti-intellectualist approach. According to their account, to know how to F is to know some proposition of the form: 'W is a way of doing F.' And this proposition is known under a '*practical* mode of presentation.' (p. 429) According to this anti-intellectualist picture, knowledge (under a practical mode) does not require an ability to cognize using representations or concepts that are mental analogues of words. A dog need not have any concept of *running* to know how to run. In knowing-how to run, the dog stands in the knowledge relation to some proposition of the form, 'W is a way of running,' but thinking of such knowledge as mediated by words or word-like constituents is utterly misleading.

This anti-intellectualist way of thinking about propositional attitudes is of central importance to Doxastic Cognitivism. Historically, cognitivists have been accused of over-intellectualizing the emotions. The intellectualist moans, '*What about the beasts and the babies?! I know that Lassie has emotions... And my baby too!*' The basic complaint is that when cognitivists posit anything like belief or judgment that P as being a constitutive element of emotion, they effectively rule out the possibility that less cognitively sophisticated creatures have emotions (like babies, or other animals that lack language.)<sup>9</sup> The intellectualist also objects by accusing cognitivists of positing a layer of conceptual structure that is completely hidden from view from the person undergoing these emotions. Objections of this kind have no force or relevance against the version of Doxastic Cognitivism that I explore.

---

<sup>9</sup> For more discussion of this intellectualist objection to cognitivism, see Deigh (1994).

### 5.3 Some Varieties of Doxastic Cognitivism

There are a variety of ways to be a Doxastic Cognitivist. While holding the above structural details fixed (e.g. positing beliefs instead of judgments, tying the relevant beliefs to states of being-V instead of states of feeling-V, and positing anti-intellectualism), let us articulate some ways to develop this view.

Let us call the strongest form of cognitivism, ‘Mad-Dog Cognitivism.’ The Mad-Dog Cognitivist holds a two-way entailment thesis between emotion ascriptions and belief ascriptions. The Mad-Dog Cognitivist will endorse strong equations like: S is angry *iff* S believes that S has been wronged or slighted, S is afraid *iff* S believes that there is some danger, etc. This two-way entailment thesis is hopelessly vulnerable to counterexamples. It is fairly easy to imagine a case where one believes the relevant proposition without being in the relevant emotional state. For example, suppose John believes by testimony that he has been slighted, it hardly seems plausible to think that this entails that John is angry. Adapting a case from Perry (1979), suppose that John is in a market and sees on the security video that a man is being pickpocketed. John believes of the victim, ‘that man is being slighted’ without realizing that he, himself, is the victim of the pickpocketing. If we also suppose that John is a very self-centered person who would never be angry if someone other than himself had been slighted, it is even easier to imagine that John is not angry. Here, we have a case where S believes S has been slighted without S being angry.<sup>10</sup>

---

<sup>10</sup> One can also make trouble for the idea that for each agent x, and each emotion, there is a belief desire pair such that x has the emotion *iff* one has the belief desire pair. If one has the belief component in the Perry (1979) case, then it will not be plausible that even if one also has the desire component, one has the relevant emotion. For discussion sympathetic to the belief-desire pair view, see Marks (1982).

Notice that the problems for Mad-Dog Cognitivism concern the direction of entailment from belief to emotion.<sup>11</sup> Perhaps a more promising form of Doxastic Cognitivism is, what I shall call, ‘Moderate Doxastic Cognitivism’ (MDC) which goes in for a one-way entailment thesis (emotion to belief). There are at least three versions of this view.

**Version 1:** For each emotional state E, there is some evaluative proposition P, such that whenever any agent is in E, that agent believes that P.<sup>12</sup>

Here is a slightly more guarded version:

**Version 2:** For each emotional state E, and each agent a, there is some evaluative proposition p such that a’s being in e entails that a believes p.<sup>13</sup>

Suppose that whenever an agent is angry, that agent believes he/she has been wronged. This idea does not secure Version 1 because the proposition *that Frank has been wronged* is not the same as the proposition that *Tim has been wronged*. But this idea does vindicate Version 2. (That is not to say Version 1 is doomed. For example, one might argue that whenever someone believes that they have been wronged, they also believe the proposition that *someone has been wronged*. This thought could be used to defend Version 1.) Here is a weaker version:

**Version 3:** For each emotional state E, there is some evaluative property F such that whenever one is in E, one has some belief involving F.<sup>14</sup>

---

<sup>11</sup> It seems implausible that for any proposition, coming to believe that proposition is accompanied by a characteristic facial expression. By contrast, Darwin observed that many emotions in humans do. See Darwin (1872).

<sup>12</sup> Solomon (1988) is clear that he denies Version 1. ‘I want to argue that the analysis of emotions as judgments must not be understood as the simple-minded claim that each emotion is the assertion of a single judgment or belief (with or without an accompanying desire).’ (p. 184)

<sup>13</sup>  $\forall x \forall y (x \text{ is an agent and } y \text{ is an emotional state}) \supset \exists p \Box (x \text{ is in } y \supset x \text{ believes } p)$

<sup>14</sup> This is a cognitivist implementation of the formal objects approach. The formal objects approach posits an intimate association between emotional kinds and evaluative kinds. The Moderate Doxastic Cognitivist explains the association by the role of those properties in evaluative beliefs. See Teroni (2007) for relevant discussion.

Suppose John feels regret on Tuesday and on Wednesday. On Tuesday, John's regret puts him in the belief relation to the proposition that the party was regrettable (where 'regrettable' is a stand-in for some evaluative property associated with regret). When John feels regret on Wednesday, John's regret puts him in the belief relation to the proposition that the golf injury was regrettable. Note that John's regret can put him the belief relation to different propositions—this observation makes no trouble for Version 3. But if we do not implicate at least one other belief that all cases of John's regret share, Version 2 is in trouble. For the purposes of our discussion, we need not favor one version of MDC over the others. But it is worth being careful to distinguish them.<sup>15</sup>

There are a few points worth emphasizing. First, Moderate Doxastic Cognitivism, as I have defined it, is non-committal on the issue of 'token identity.' That is to say, it is non-committal about whether the event of having an emotion is identical to the event of having the belief whose existence is entailed by the emotion. Consider an analogy. It is plausible that punching Jones entails making contact with Jones. On a particular occasion, I punch Jones. The entailment does not straightforwardly settle whether, on that occasion, the event of punching is identical to the event of making contact. I wish to remain neutral on these tricky issues of event identity. For that reason, we should not think of MDC as the view that emotion events or emotional states are identical to belief-events or belief-states. Similarly, except when speaking loosely, we should be wary of the slogan that emotions *are* beliefs.

---

<sup>15</sup> One issue relevant to Versions 1 and 2: Are we to think that for every degree of anger *d*, and every agent *x*, there is a different proposition corresponding to the state of being angry to degree *d*? Note that the one-way-entailment thesis does not require this.

Second, in positing a one-way entailment, the Moderate Doxastic Cognitivist does not preclude that for some emotional V's, Ving entails that an agent has certain desires.<sup>16</sup> Consider an analogy. Let a 'foon' be a utensil that is a fork at one end and a spoon at the other. (To be carefully distinguished from a *spork*, which is an amalgam of a fork and a spoon at the same end.) If we are counting the number of spoons in the world, we will include the foons. If we are counting the number of forks in the world, we will also include the foons. (We might want to remain neutral on the contentious issue of whether or not we ought to include the sporks in our count of forks too.) Moderate Doxastic Cognitivism is compatible with the idea that some emotions are like foons. For such an emotion, having that emotion requires having characteristics that suffice for a certain belief state and in addition requires having characteristics that suffice for a certain desire state. For example, perhaps fear entails both a belief concerning danger and a desire to avoid danger.

Let us now consider another view which I will call 'Guise-Cognitivism.' The Guise-Cognitivist posits an emotion-to-belief entailment but makes a more specific claim than the Moderate. For Version 1, the Guise-Cognitivist says for each kind of emotional state, there is some guise (or set of guises), and some proposition, such that to be in that particular kind of emotional state is to believe a particular proposition under that guise. (There are guise-varieties of Version 2 and Version 3, but we need not explore these views further here.) The guise theorist does not accept Mad-Dog Cognitivism, since believing a proposition may not entail *any* emotion when believed under an unsuitable guise. The guise theorist does however accept the one-way entailment view, since emotion entails a belief in a proposition.

---

<sup>16</sup> Solomon (2003, p.106) writes 'one might analyze the various emotions as judgmental structures enclosing a core desire which is both their motivation and their 'conatus.'

What exactly is meant by ‘guise?’ The Fregean-literature is suggestive of a model that is useful here: the ‘modes of presentation’ model.<sup>17</sup> Here, the main idea is that one way to believe that Venus is bright is under the ‘Hesperus-guise’ (as one does when one believes that Hesperus is bright), and another way to believe that Venus is bright is under the ‘Phosphorous-guise’ (as one does when one believes that Phosphorous is bright). For the guise-cognitivist, an emotion is not an ordinary way of having a belief that P. Rather, having an emotion is a *special way* of having a belief that P. Thus, John’s guilt is tantamount to John believing that he has done something wrong under the guise of his guilt.

To get a feel for the Guise-Cognitivist approach, it may be useful to return to Stanley and Williamson’s (2001) analogous account of know-how. On this model, knowing how to F is knowing that a certain activity is a way of Fing under a practical-*guise*. Thus, knowing how to ride a bike is knowing (in a practical/skillful way) a proposition that says of *some way* that it is a way of riding a bike. Contrast this kind of knowledge with the knowledge-how that one might have if they merely had an instruction manual for bike-riding. Upon learning the conjunction of the propositions provided by such a manual, one gains intellectual-know-how (knowing how to F under an intellectual-guise) inasmuch as one possesses the linguistic knowledge sufficient for knowing how to ride a bike. But intellectual know-how is not sufficient for knowing how to ride a bike in the natural sense—which is knowing how to ride a bike under the a practical-guise. By analogy, the Guise-Cognitivist would say that John’s guilt is a special way of believing that he has done something wrong under the ‘guilt-guise’. In other words, John’s guilt is special way of standing in the belief relation to the proposition that *he has done something wrong*. This is one

---

<sup>17</sup> See Frege (1892).

way to illuminate the contrast between John and a guilt-free psychopath (who could only know that he has done something wrong under a non-affective, intellectual-guise).

#### **5.4 Inconclusive Arguments Against Moderate Doxastic Cognitivism**

Let us now turn our attention to objections to Doxastic Cognitivism, with the Moderate view in mind. Let us also remain neutral about the guise-implementation of Moderate Doxastic Cognitivism.<sup>18</sup> Many of the arguments in the literature appear as arguments against the view that emotions entail evaluative judgments. Where necessary, I have adjusted the target of these objections to see if Doxastic Cognitivism will fare any better.<sup>19</sup>

##### **Superficial Linguistic Argument 1**

One might reject cognitivism on the grounds of a very superficial linguistic argument: Although we talk about beliefs as being the sorts of things which can be either true or false, we do not usually talk about the emotions as though they are true or false.<sup>20</sup> If we ever say that John's anger was true, we would likely mean to say that his anger is genuine.<sup>21</sup> In those special cases where we might call an emotional state true (e.g. John's love, John's happiness, John's remorse, etc.), it seems we are really just saying something about the sincerity or genuineness of John's emotion. One might thus reject cognitivism on the grounds that the different ways we talk about the beliefs and emotions is at least suggestive that they are very different kinds of things.

---

<sup>18</sup> One reason for remaining neutral is that the task of evaluating Guise-Cognitivism would require a careful account of what can and cannot count as a 'mode of presentation' – without any constraints on this, it is difficult to get a clear fix on the problems this particular view might have. For our purposes, we need not explore this view any further here.

<sup>19</sup> For more extensive discussion of standard non-cognitivist objections, see: De Sousa (1987); Greenspan (1988); Deigh (1994); Griffiths (1997); Goldie (2000); DeLancey (2001); D'Arms and Jacobson (2003); Roberts (2003); Robinson (2004); Prinz (2004); and Scarantino (2010).

<sup>20</sup> See Astington (1993, p. 69).

<sup>21</sup> There is also a use of 'false anger' but it has little to do with the use of 'falsehood' applied to belief.

And yet cognitivism predicts that emotional states involve contentful beliefs that can be true or false and if this is right, then emotions do involve truth/falsity.

The superficial linguistic argument overlooks the fact that there are plenty of other kinds of states that are obviously like belief but for which the language of truth or falsity is inappropriate.<sup>22</sup> Consider the proposition that *she is confident that she will win*. It is plausible that her confidence suffices for belief but it is very odd to say that her confidence is ‘true’ unless we mean to say that it is genuine. (It is similarly odd to say that her confidence is ‘false’ unless we mean to say that it is somehow not genuine.) Note that even if we managed to convince ourselves that being confident that P is not sufficient for having a belief that P, we still have the option of saying that being angry is a way of being confident of a certain evaluative proposition (e.g. that one has been wronged)—and even if this does not amount to saying that anger involves belief, it is still in the spirit of a cognitivist view. In this connection, it is also worth considering the notion of ‘acceptance’ which may differ in subtle ways to belief. It is just as strange to apply ‘true’ or ‘false’ to an agent’s acceptance (whereby S accepts that P and the acceptance object is some evaluative proposition). The superficial linguistic argument would fail again if we supposed (instead of cognitivism) that various emotions constitutively involve acceptances of certain evaluative propositions (an idea that is very much in the spirit of cognitivism).<sup>23</sup>

Knowledge and certainty can deliver the final blow to the superficial linguistic argument. For knowing clearly suffices for belief and yet the language of truth and falsity remains similarly problematic. For example, consider the claim, ‘She knows that P’. Factivity precludes us from saying that her knowledge is false. But interestingly, if we were to say that her knowledge is true,

---

<sup>22</sup> Note also that entailing the presence of something true or false is not automatically sufficient for being true or false.

<sup>23</sup> See Helm (2001) for the view that emotions involve something like assent to an evaluative proposition.

we could only mean to convey something about genuineness. Admittedly, one might say that this reading is blocked by considerations of redundancy (knowledge is already factive so why bother characterizing her knowledge as ‘true’?) Consider, ‘She is certain that P’. It is very strange to say that her certainty is true *except* as a means to convey genuineness but this cannot be explained by appealing to redundancy. ‘She is certain that P’ is not factive (it does not entail truth) even though ‘It is certain that P’ is factive. In sum, the superficial linguistic argument is a complete failure.

### **Superficial Linguistic Argument 2**

There is another superficial linguistic argument against cognitivism.<sup>24</sup> Roughly, the idea is that a verb (V) expresses a real propositional attitude only if, ‘S V’s the proposition that P’ is a sensible thing to say of an agent (S). Perhaps, more precisely, a verb expresses a real propositional attitude only if the inference from ‘S V’s that P’ to ‘S V’s the proposition that P’ is valid. For example, the inference from, ‘Jones knows that he has a train ticket’ to ‘Jones knows the proposition that he has a train ticket’ is arguably valid. Thus, it is arguable that knowledge passes the test and is a genuine propositional attitude. Emotions will invariably fail this test. It is obvious that the inference from, ‘Jones is afraid that the rabbit will bite’ to ‘Jones is afraid of the proposition that the rabbit will bite’ is terrible. I am skeptical that this is a good test for whether an emotional verb expresses a propositional attitude. When we apply this test more generally, it predicts that a range of attitudes which are standardly thought to be propositional attitudes are not propositional attitudes after all. Setting emotions aside, this test predicts that ‘suspects that P’, ‘discovers that P’, and ‘thinks that P’, etc. are not genuine propositional attitudes. I take it that this is a very unsatisfying result.

---

<sup>24</sup> One will find this objection in Deigh (1994, p. 846).

## The Problem of Recalcitrant Emotions

Recalcitrant emotions are thought to present another problem for cognitivism. A recalcitrant emotion is one that persists ‘despite the agent’s making a judgment that is in tension with it...’ (D’Arms and Jacobson, 2003, p. 129) The basic idea here is that one can have an emotion like fear while holding a conflicting evaluative belief.<sup>25</sup> Thus, one might have a persisting fear (which, on my view, entails some kind of belief along the lines that one is in danger) despite knowing that one is not in any danger at all. The problem of recalcitrant emotions is not a conclusive argument against cognitivism. There are a few obvious ways for the Doxastic Cognitivist to respond.

First, it is important to recall that one can believe something under one guise while having a belief in its negation under another.<sup>26</sup> John can believe that Hesperus is big and believe that Phosphorous is not big without thereby being strange or irrational. In a case where John does not know that Frank is Jack the ripper, John can believe that Frank is innocent under the Frank-guise and believe that Frank is guilty under the Jack the Ripper guise. It is fairly commonplace to believe that P under one guise and believe that not P under a different guise.

Of course, the guise response may not adequate by itself. Once John knows that Frank is Jack the ripper, there is overwhelming pressure to adjust one of the earlier attitudes, on pain of serious irrationality. If one thinks of the emotions as guises (as the Guise-Cognitivist does) then one would expect that either the emotion or the conflicting judgment would be retracted once they are acknowledged to be presentations of conflicting propositions. For example, consider a case of recalcitrant-guilt. Suppose S is guilty about some action, F, while knowing that there is nothing wrong with doing F, and while knowing that S’s guilt is out of place if F is not wrong.

---

<sup>25</sup> See Michael Brady (2009) for a more detailed account of recalcitrant emotions.

<sup>26</sup> See for example, Crimmins and Perry (1989).

The argument against Doxastic Cognitivism would have to be to something along the following lines: (i) S is in a situation that is pretty ordinary for all agents. (ii) This kind of situation would not be ordinary on the assumption of Doxastic Cognitivism. Therefore, Doxastic Cognitivism is wrong.

The problem with this argument is that even when we focus on beliefs, the kind of conflict we are talking about is not so strange.<sup>27</sup> Someone that was brought up Jewish but no longer practices that religion might believe that there is nothing wrong with eating pork, and yet think that they are doing something wrong when eating bacon. Someone might know that their partner is dead but still wake up in the morning expecting to turn over and see them in the bed. Theorizing about these situations is tricky but in the cases above, there is perfectly good sense in which one thinks that something is so while also thinking that it is not. The Doxastic Cognitivist can reply to the objection of recalcitrant emotions by assimilating these recalcitrant-emotion cases (like the guilt-case above) to cases like these. The key idea is that there is often a real mismatch between certain beliefs and one's considered view about the way things really are.

In a similar vein, one might object to Doxastic Cognitivism by arguing that since recalcitrant emotions are an ordinary occurrence, Doxastic Cognitivism makes agents seem more irrational than we should think that they are. Brady (2009, p. 414) recommends against positing conflicting states (conscious or not) by appealing to Greenspan's idea of 'logical charity' which suggests against ascribing complex evaluative profiles to agents if, 'We need some special reason ... for attributing to him an unacknowledged judgment in conflict with those he acknowledges.'

Greenspan (1988, p. 18)

---

<sup>27</sup> As Helm writes, 'conflicts between emotions and judgements do not verge on incoherence, for they are readily intelligible and happen all too often.' (2001, p. 42)

Regardless of whether Cognitivism is true, we are usually quite happy to call emotions ‘irrational’ sometimes. Suppose John knows that he is not in danger but feels fear. Of course, we would say that John is being irrational. Denying Cognitivism does not protect agents from irrationality. Whether or not we think that fear entails belief, we would treat John’s fear as irrational. It is not clear that Doxastic Cognitivism makes recalcitrant emotions an excessive sign of irrationality.

What are we to make of Greenspan’s idea of ‘logical charity’? Granting that when a person believes they did nothing wrong, we need some kind of reason to claim that the person believes that they did (as we might in the case of guilt). But I am not sure what ‘special’ is supposed to add. In general, when we believe that someone has a belief, we ought to have evidence that they do. Doxastic Cognitivism (as I have developed it) encourages the idea that when an agent is angry, we have an excellent reason to think the agent believes that they have been wronged or slighted. (In passing, note that there is no deep worry here that the agent may not explicitly acknowledge that they believe that they have done something wrong – for there is nothing unusual about this. We often have very compelling grounds for attributing belief to an agent even though the agent will not explicitly claim that they have it.)

### **The Specification Challenge**

One attraction of the cognitivist picture was taken to be that it could potentially provide a way of individuating different kinds of emotions (like fear versus anger, or anger versus sadness) by appealing to their different propositional contents. But this supposed ‘upshot’ was turned as a challenge against cognitivism: stating what the proposition is. For any given emotion (like say,

anger), it is very difficult to specify the proposition that is believed when one is angry.<sup>28 29</sup> Here is a bad way to run the challenge: ‘We can recognize what emotions we are having without needing to first recognize that we have a particular attitude to a proposition, so we should not think that having an attitude to a particular proposition is essential to having that emotion.’ This is a terrible thought. Being H<sub>2</sub>O, as Kripke (1980) reminded us, is essential to being water. That said, it took centuries of theoretical inquiry to figure this out. English peasants did not need to know chemistry to know that some puddle was a bit of dirty water. (See Aristotle on the order of being and the order of knowing).<sup>30</sup> A better way to run the challenge is by running it toward the theorist: In the case of chemistry, we have a pretty good idea how to figure out the exact chemical composition of a substance but in psychological cases, it is scarcely imaginable that we can find out the precise propositional composition of an emotion.

There are at least two promising ways the cognitivist might respond to this specification challenge. One way is to appeal to vagueness and the other is to appeal to the idea that the relevant propositions are *sui generis*.

The first way the Doxastic Cognitivist can try to answer the specification challenge is by appealing to vagueness. For example, if anger involves (at least, in part) believing a particular proposition, the question of what exactly that proposition is, is a vague matter. Just as we cannot give precise answers to questions about borderline cases (e.g. whether certain people are bald or whether, on a certain day, twilight began at t<sub>1</sub>, or t<sub>2</sub>, or t<sub>3</sub>, etc.), there is no obvious way that we

---

<sup>28</sup> This problem is especially vivid if we consider graded-emotional states (e.g. being surprised that P versus being very surprised that P, or being a little afraid that P versus being terrified that P, or being angry that P versus being livid that P, etc.

<sup>29</sup> The specification challenge differs in accordance with the version of MDC one goes for. For example, in the context of Version 3, the challenge will be to specify the evaluative property is associated with each emotional state.

<sup>30</sup> See *Categories and De Interpretatione*. Aristotle (1963)

should expect to answer the question: *What proposition is believed when one is angry, afraid, terrified, etc.?*

Appealing to vagueness does not mean that we have nothing to say about the content of emotion. We are merely limited in what we can say. Consider an analogy. Suppose Frank uploads a low-resolution photo of himself to his online dating profile. The photo is of such low-quality, it seems to conjure a representation of a tall, dark, and lithe man. Although the image is not detailed enough to inform his prospective date about the little aesthetic details (e.g. whether he has nice teeth or youthful skin) the photo still presents a rough approximation of Frank's looks. Even if the prospective date cannot tell whether the lighting of the photo is accurate, the prospective date can still gather that Frank is tall.

The second way of responding to the specification challenge is to say that we should not assume that every proposition can be easily expressed using the resources of ordinary language. Rather, it is better to think of the proposition that the Doxastic Cognitivist has in mind as *sui generis*, a thing of its own kind, which cannot be informatively expressed using intellectualist resources (e.g. sentences).<sup>31</sup> Note that this move is not incompatible with the vagueness approach. Perhaps, when one is angry, it is vague which *sui generis* proposition one believes.

*Further Worries.* Some arguments against cognitivism are not of serious concern for the Doxastic Cognitivist because they either specifically target an intellectualist, judgment version of cognitivism or because they target Mad-Dog Cognitivism. One will find a number of objections to intellectualist-cognitivism in Andrea Scarantino's, 'Insights and Blindspots of the Cognitivist Theory of Emotions' (2010). For example, Scarantino's observation that emotions entail motivational states is not even *prima facie* a worry against the one-directional Moderate

---

<sup>31</sup> Of course, one could use 'Jones' to refer to the proposition that is associated with anger. This would count as using natural language to express what the proposition is without being informative.

Doxastic Cognitivism. It would be a worry for Mad-Dog Cognitivism (at least if we assume that no belief entails motivational states). Nor is his observation that emotions have a distinctive phenomenology while judgment does not guarantee a phenomenology, since, even granting that belief similarly does not guarantee a phenomenology, this is no problem for Moderate Doxastic Cognitivism. Consider further his thesis that *blindfright*—the phenomenon of fear without awareness of the triggering stimuli for fear—is a problem for cognitivism. It is not plausible that awareness of the source of one's belief is a condition on belief. Hence *blindfright* poses zero threat to Doxastic Cognitivism.<sup>32</sup>

Against Doxastic Cognitivism, the objection that 'emotions can be elicited unconsciously and without deliberation, whereas judgments are deliberate and conscious' is also a non-starter.<sup>33</sup> When we adapt the target of the argument from judgment-based cognitivism to Doxastic Cognitivism, it is perfectly obvious that beliefs can be elicited unconsciously and without deliberation.

The discussion thus far has been clear that the Doxastic Cognitivist is a non-intellectualist about belief, and is not at all threatened by the fact that animals and young children have emotions. Thus, when Scarantino (2010) expresses the concern that such creatures, 'do not have the language mastery required for giving or withholding consent of the impression that something is the case' (p. 743) the Doxastic Cognitivist can dismiss the objection for its irrelevance. Scarantino expects the traditional cognitivist to try to stave off refutation by opting

---

<sup>32</sup> Scarantino (2010) provides a good discussion of the weaknesses of intellectualist-cognitivism.

<sup>33</sup> For example, Brady (2009) criticizes judgment-based cognitivism on the grounds that 'subjects who suffer from recalcitrant emotions do not consciously assent to the judgement that is supposedly constitutive of their emotion...'. (p. 414) Later, Brady writes that to believe that X is F, one has to 'assent to the appearance' (p. 416) but if 'assent' means anything like *conscious assent*, Brady is wrong. (Of course, Brady is making an uncontentious point if 'assent' merely means believe.)

for an extended notion of judgment, one that is so ‘elastic’ as to make cognitivism devoid of substantive content. Instead, I recommend that the cognitivist stave off refutation by adopting Moderate Doxastic Cognitivism. It is hard to see that this move would somehow make cognitivism devoid of interesting content, yet it escapes the bulk of his critique by simultaneously switching from a judgment-oriented view to a belief-oriented view, and by distancing itself from Mad-Dog Cognitivism.

### 5.5 A Positive Case for Doxastic Cognitivism

To make a positive case for Moderate Doxastic Cognitivism, we should ask whether the emotions play a functional role that is characteristic of the functional role beliefs play. To answer this question, it is helpful to separate two aspects of the functional role beliefs play—their input role (e.g. what causes beliefs, what properties beliefs track/are correlated with, etc.) and their output role (e.g. the role of beliefs in producing action).

We might say that at the input end, beliefs serve to ‘indicate’ that certain states of affairs obtain.<sup>34</sup> For example, in an optimal case where an agent has a belief about their environment (e.g. that the berry bush is poisonous), we think of the agent as standing in the knowledge relation to their environment. Here, we might draw on a safety conception of knowledge where the basic idea is that a belief that P is safe just in case the belief could not easily have occurred without P occurring.<sup>35</sup> If emotional states are not the kind of states that can be safely correlated with evaluative states of affairs (in the environment), then we ought to be pessimistic that emotions (like fear, guilt, anger, ect.) entail beliefs. On the other hand, if emotions *can* safely

---

<sup>34</sup> Here, I am borrowing the language of Robert Stalnaker (1984) but we need not get into a lengthy discussion of what it is for an ‘indication’ to happen (and in particular, we need not delve into Stalnaker’s idea that to indicate P is to be in a state that under optimal conditions occurs because of P.)

<sup>35</sup> See chapter three for a detailed discussion of safety.

track evaluative states of affairs (and I take it that they can and do) then this is pretty compelling evidence that emotions are well-suited to play a similar functional (input) role that belief plays.

Suppose a street performing clown intentionally runs over my feet on his unicycle. A state of anger is induced. It is very plausible that in such a case, my state of anger is safely correlated with the clown wronging me. That is to say, I could not easily have been in that very state without the clown wronging me. Of course, I could easily have been angry without the clown wronging me—for even if he hadn't, I might still be angry that Frank once had an affair with my mother. But it seems reasonable to ask whether the state of anger that the clown produced could easily have occurred without his wronging me. In many situations, it is natural to think that the relevant emotional state could *not* easily have occurred without the clown intentionally running over my feet. (Note that it is a further question whether it could easily have happened that the clown wronged me without causing my anger. But that is not relevant to the safety idea: Safely believing P means that there is no relevant danger of being mistaken about P, not that there is no relevant danger of my being ignorant that P.)

Similarly, it seems right to think that there are fear states that safely track the presence of danger. When someone runs into the bank with a gun and yells, 'Everyone get down on the floor! This is a stick-up!' I am overcome with fear. The state of fear that is produced could not easily have happened in the absence of the danger. In sum, it seems that emotional states are fit to safely track evaluative states of affairs. At the input end of things, emotions seem to play a knowledge-role, and anything fit to play a knowledge-role at the input end is not going to be shown to be a non-belief on the basis of input considerations alone.

One payoff that comes from taking emotions to indicate evaluative states of affairs is that it offers a promising way to account for the truth of emotional state attributions which take a

noun phrase (NP) complement. For example, suppose we ask, ‘What makes it true that Frank regrets *the party*?’ Assuming that we are thinking of cases where Frank has personal reasons for the regret, we might suggest that Frank regrets the party when *the party* figures in his reasons for regret. But this is not an acceptable theory. For example, suppose I organize a party and then find out a second, more awesome party (the ‘Super-party’) is going on at the same time as my event. I regret my party on account of the fact my party is making me miss the Super-party. In this case, Super-party figures in my reasons for regret but I do not regret the Super-party. The indication approach has a potentially promising explanation. Suppose that any episode of regret indicates some proposition of the form *X is regrettable* (where ‘regrettable’ is a stand in for whatever normative property is associated with regret). In this setting, it is plausible that to regret some object/event *x*, is for *x* to figure in a content of the form *x is regrettable* that is indicated by some state of regret.

Of course, the argument above is not a decisive argument for Doxastic Cognitivism. For one might think that emotions indicate evaluative propositions and use the indication approach above to argue that the evaluative propositions indicated by emotions are a basis for explaining the emotions without committing oneself to the thesis that emotions entail beliefs.<sup>36</sup> Let us turn to the question of whether emotions satisfy the functional requirements of belief on the output end.

The discussion above suggests that the emotions can safely track evaluative states of affairs, and this is pretty good evidence that they are not disqualified at the input end from being the ground of true belief ascriptions (since they are not generally disqualified at the input end

---

<sup>36</sup> I realize that someone might think that emotional relations to objects described by such constructions as ‘regretting that party’ cannot be reduced to or fully explained by attitudinal relations to propositions. For relevant discussion see Deigh (1994, sections iv-vi.) It is beyond the scope of this essay to pursue this thread further.

from being the ground of true knowledge ascriptions). Yet that is hardly a compelling case (by itself) for Doxastic Cognitivism. As Stalnaker (1984) emphasizes, the idea that a state plays some indication-role is hardly a sufficient basis for that state count as a belief. For example, Stalnaker writes,

... if a bald head is shiny enough to reflect some features of the environment, then the states of that head might be described in terms of a kind of indicator – in terms of a relation between the person having the head and a proposition. But no one would be tempted to call such states belief states. (1984, p. 18)

The states of the bald person's scalp do not play a belief-like role with respect to reasoning and action. We should not think that the states of the bald person's scalp are the right kinds of things to count as belief (even though they may very well be safely correlated with states of affairs in the world).

Consider the fact that certain neurons in the brain are responsive to skin temperature and play crucial roles in maintaining homeostasis (e.g. registering body temperature, triggering appropriate responses in muscles and sweat glands.) (Zhao et al., 2017)<sup>37</sup> It is clear that these neurons (or states of those neurons), in a very natural sense, register and indicate states of skin temperature. There is even a very natural sense in which we can say these neurons *represent* skin temperature. But we should not think that having these neurons in the hypothalamus is tantamount to having beliefs about our skin temperature (even if those states have *the function* of being reliably connected to states of the world).<sup>38</sup>

---

<sup>37</sup> Certain neurons in the preoptic area and also the dorsomedial hypothalamus are responsive to skin temperature. (See Zhao et al. 2017)

<sup>38</sup> Some theorists (e.g. Prinz, 2004) think that this idea of function is of central importance to representation.

For our purposes, we need not worry about whether emotions are representational (it is fairly obvious that they are in some sense). We should not think that the question of whether emotion entails belief will turn on whether emotions represent—representation comes cheap (as is clear from the example of hypothalamic neurons above).<sup>39</sup>

The merits of Doxastic Cognitivism will hinge on whether emotions play a primary functional role that we presume beliefs to play: namely, interacting with other states (e.g. belief and desires) to produce action.<sup>40</sup> Consideration of the output-role of emotion will allow us to evaluate the competing merits of alternative views (e.g. views that say emotions are impressions or construals or Stoic *phantasma* that are less committal than belief).

A great deal of our knowledge is not directly manifested in conscious reasoning. As I walk around, I know various things about the relative positions of my limbs though I am not readily able to articulate all that knowledge in explicit conscious reasoning using sentences or sentence-like vehicles. Much that same can be said of vast swathes of procedural knowledge when we assume Stanley and Williamson's (2001) thesis that this involves propositional knowledge under practical modes of presentation. The fact that emotions are not components of conscious reasoning does not settle whether they are belief-like.

What is a reasonable output-related test for whether a state is belief-like? The prevalent idea in the literature about the output role of beliefs has to do with the way they combine with desires to produce action. We can find this representative idea described by Stalnaker below.

---

<sup>39</sup> Jesse Prinz (2004, Ch. 3) devotes himself to the task of defending a Jamesian-thesis that emotions represent bodily states by using a two-fold test of whether (a) they are reliably triggered by bodily states and (b) whether they have the function of doing so.

<sup>40</sup> Note that if the presumed functional role of beliefs is that they figure in as explicit premises in deductive reasoning, then emotions will not fare well. Emotions do not seem to figure as premises in explicit deductive reasoning. But we should not think that this discredits emotions from doxastic status especially if judgments (not beliefs) are the fodder of linear explicit conscious reasoning. After all, I am not trying to defend the view that emotions are judgments.

To desire that P is to be disposed to act in ways that would tend to bring it about that P in a world in which one's beliefs, whatever they are, were true. To believe that P is to be disposed to act in ways that would tend to satisfy one's desires, whatever they are, in a world in which P (together with one's other beliefs) were true. (1984, p. 15)

This test does not give us a simple way of testing for the presence of belief unless we have already established which desires are in place. (An additional complication arises when we consider guises. If I believe Jack is nearby under the Frank-guise and I want to kiss Frank under the Jack-guise, then my belief will not combine with my desire to kiss Jack to motivate me to kiss Frank.)<sup>41</sup> But this test does give us a rough and ready way to probe the issue by allowing us to ask, 'Are states of regret, fear, anger, etc. the kind of states that can combine with desires in a way beliefs in evaluative propositions do?' If emotions are not the kinds of states that can combine with other beliefs and desires to motivate action, then the Doxastic Cognitivist is in trouble. But there is a very good case to be made that emotions do so-combine.

Suppose I fear the snake. On the hypothesis that fear entails that I believe the snake to be a source of danger, it seems pretty obvious that my fear coupled with my desire to avoid danger, would incline me to avoid the snake. It certainly seems true that fear coupled with a desire to avoid danger (setting aside the subtleties of the snake being presented under multiple guises) would automatically incline one to avoid the snake (also setting aside that there may be competing desires). This kind of interaction is exactly what one would expect from the combination of Doxastic Cognitivism coupled with the standard conception of the role of belief in generating action.

---

<sup>41</sup> See Stalnaker (1984) for his own extensive account of this issue.

Similarly, consider a version of Doxastic Cognitivism which says that guilt about a past action entails a belief that the past action was wrong. This would predict that such guilt in combination with a desire to confess past wrongdoings that one hadn't already confessed, and knowledge that one had not yet confessed yet, would incline one to confession. And this seems to be exactly the right prediction about the role that the guilt would play. (Consider a contrast. A 'seeming of wrongdoing' or an 'appearance of wrongdoing' might motivate one to seek out reasons/evidence that bears on the accuracy of that appearance.<sup>42</sup> Such an appearance would not, by itself (even in combination with a desire to confess misdeeds) induce a confession. On the other hand, guilt *does* combine with such a desire in a way that the impression of guilt would not. So much worse for the view that guilt is a mere impression of wrongdoing.) Suppose John has a choice between going to party-A or party-B, he prefers A to B if Tim is there and party-B otherwise. John is optimistic that Tim will be at party-A. John's optimism in combination with his preference, pretty much ensures that he will go to party-A. Here, optimism seems to play the kind of role that you would expect it to play if you thought optimism involved a belief that something is likely.<sup>43</sup> In short, when asking, 'When does an episode of emotion underwrite the truth of a certain belief attribution?' we should look at whether the emotional state is functionally an adequate basis for belief attribution. And there is at least some basis for being optimistic here.

There is a further way that emotions are belief-like. Beliefs have offline analogues that we deploy under supposition.<sup>44</sup> When we suppose that P, we offline believe the very sorts of

---

<sup>42</sup> See Brady (2013) for an account of the emotions which roughly says that emotions are merely motivational resources which prompt the search for reasons or evidence that bear on their accuracy. See chapter 3 for a critical examination of Brady's proposal.

<sup>43</sup> Someone might object that if the stakes are high, John will delay his choice and instead seek further confirmation that Tim is at the party. That is true enough but it by no means disconfirms that view that optimism involves a belief that something is likely. After all, if John believes that Tim is likely to be there but the stakes are high, John might double check and then make a choice.

<sup>44</sup> See chapter four.

things that we believe when we learn P. It is wrong to think that all mental states have offline versions. For example, it is not right to think there is such a thing as an offline-toothache that we have under suppositions. Of course, I may be so wired that supposing certain things causes me to have a toothache but this is not an offline-toothache, it is an actual toothache! And of course, I can imagine/suppose/talk about having toothaches, but these are not offline-toothaches either. In this respect, emotions are much more like beliefs than pains. Under suppositions, we can have offline-anger, offline-delight, etc. Of course, their capacity to exist in offline-forms does not suffice to establish that emotions in their online-form are beliefs but it is conducive to the Doxastic picture. Overall, the combination of input and output considerations given above seem to offer a pretty good *prima facie* case for Doxastic Cognitivism.

## 5.6 Concluding Remarks

In this essay, I have presented a novel case for a cognitivism that is both doxastic and anti-intellectualist. We have seen that many of the standard objections to judgment-based cognitivism look far weaker when transposed to Doxastic Cognitivism. We have also seen that a pretty strong case can be made in favor of Doxastic Cognitivism. Of course, there is plenty of further work to be done in critically evaluating Doxastic Cognitivism and exploring its alternatives. Against the pull of the literature, the idea that emotions carry doxastic commitments to evaluative propositions looks surprisingly robust.



## BIBLIOGRAPHY

- Aristotle. (1963). *Aristotle: Categories and De Interpretatione*. (J. L. Ackrill, Trans.) Oxford: Clarendon Press.
- Anderson, A. K., & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature*, 411(6835), 305-309.
- Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford: Oxford University Press.
- Astington, J. (1993). *The Child's Discovery of the Mind*. Cambridge, MA: Harvard University Press.
- Blumberg, K. & Holguin, B.. (2018). 'Embedded Attitudes'. *Manuscript Draft*.
- Bodenhausen, G. V. (1993). Emotion, arousal, and stereotypic judgments: A heuristic model of affect and stereotyping. In D. M. Mackie, & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 13-37). San Diego: Academic Press.
- Bodenhausen, G. V., Gabriel, S., & Lineberger, M. (2000). Sadness and susceptibility to judgmental bias: The case of anchoring. *Psychological Science*, 11(4), 320-323.
- Brady, M.S. (2009). The irrationality of recalcitrant emotions. *Philosophical Studies*. 145. 413-430.
- Brady, M.S. (2011). Emotions, Perceptions, and Reasons. In C. Bagnoli (Ed.), *Morality and the Emotions*. Oxford: Oxford University Press.
- Brady, M.S. (2013). *Emotional insight: The epistemic role of emotional experience*. Oxford: Oxford University Press.
- Buchanan, T. W. (2007). Retrieval of Emotional Memories. *Psychological Bulletin*, 133(5), 761–779.
- Chalmers, D., Hajek, A. (2007). Ramsey + Moore = God. *Analysis*, 67.72.
- Crane, T. (2001). Intentional Objects. *Ratio* 14 (4):298-317.
- Crane, T. (2003). The Intentional Structure of Consciousness. In Quentin Smith & Aleksandar Jokic (Eds.), *Consciousness: New Philosophical Perspectives*. Oxford, UK: Oxford University Press. pp. 33-56.
- Crimmins, M. and Perry, J. (1989). 'The prince and the phone booth: Reporting puzzling beliefs' *Journal of Philosophy* 86: 685-711.
- Dancy, J. (2000). *Practical Reality*. Oxford University Press.
- D'arms, J. and Jacobson, D. (2003), 'The Significance of Recalcitrant Emotion (or antiquasijudgementalism)', in Hatzimoysis, A. (Ed.), *Philosophy and the Emotions: Royal Institute of Philosophy Supplement*: 52 (Cambridge: Cambridge University Press), pp. 127–146.

- Darwin, C. (1872/1998). *The Expression of the Emotions in Man and Animals*. London: Harper Collins.
- Davis, M., & Whalen, P. J. (2001). The amygdala: Vigilance and emotion. *Molecular Psychiatry*, 6, 13–34.
- Dietz, C. H. (2016). Are all reasons causes? *Philosophical Studies*, 173(5), 1179–1190.
- Dietz, C. H. (2017). Reasons and factive emotions. *Philosophical Studies*, (7):1-11.
- Deigh, J. (1994). Cognitivism in the Theory of Emotions. . *Ethics*, 104(4), 824-854.
- De Gelder, B., Vroomen, J., Pourtois, G., & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *NeuroReport: For Rapid Communication of Neuroscience Research*, 10(18), 3759-3763.
- Dearborn, G. V. N. (1918). *The psychology of clothing*. Princeton, N.J: Psychological Review Company.
- Descartes, R., (Ed.) Cottingham, J., Stoothoff, R., & Murdoch, D. (1984). *The philosophical writings of Descartes*. Cambridge. Cambridgeshire: Cambridge University Press.
- De Sousa, R.B. (1987) *The Rationality of Emotion*. MIT Press.
- DeSteno, D., Li, Y., Dickens, L., & Lerner, J. S. (2014). Gratitude: A Tool for Reducing Economic Impatience. *Psychological Science*, 25(6), 1262–1267.
- Dolan, R.J. (2002). Emotion, Cognition, and Behavior. *Science*, 298 (5596)., 1191-1194.
- Drucker, D. (2017). Policy Externalism . *Philosophy and Phenomenological Research*.
- Elgin, C. Z. (2008). Emotion and Understanding . In U. D. G. Brun (Ed.), *Epistemology and Emotions*.
- Ferneyhough, E. K. (2013). Anxiety modulates the effects of emotion and attention on early vision. *Cognition & Emotion*, 27(1), 166–176 .
- Fodor, J. (1978). Propositional Attitudes. *The Monist*, 61(4), 501-523.
- Frege, G. (1879/1997). Begriffsschrift. In Beaney (Ed.). *The Frege Reader*. Oxford: Basil Blackwell.
- Frege, G. (1892/1960). 'On Sense and Reference,' In *Translations from the Philosophical Writings of Gottlob Frege*. (P. G. Black, Trans.) Oxford: Basil Blackwell.
- Frege, G. (1906/1997). 'A Brief Survey of My Logical Doctrines'. In G. Frege, & Beaney (Ed.). *The Frege Reader*. Oxford: Basil Blackwell.
- French, C. (2012). Does propositional seeing entail propositional knowledge? *Theoria*, 78(2), 115–127.
- Goldie, P. (2000). *The Emotions: A Philosophical Exploration*. Oxford University Press.
- Goldie, P. (2004). Emotion, feeling, and knowledge of the world. In Robert C. Solomon (Ed.), *Thinking About Feeling: Contemporary Philosophers on Emotions*. Oxford University Press.
- Gordon, R.M. (1987). *The Structure of Emotions: Investigations in Cognitive Philosophy*. Cambridge University Press.
- Greenspan, P. (1988). *Emotions and Reasons: An Inquiry into Emotional Justification*. New York: Routledge, Chapman and Hall.

- Greco, J. (2007). Worries about Pritchard's Safety. *Synthese*, 158: 299.
- Griffiths, P. E. (1997). *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- Grice, H. P. (2001). *Aspects of reason*. Oxford: Clarendon Press.
- Hawthorne, J., Magidor, O. (forthcoming). Reasons and evidence. In D. Star (Ed.), *The Oxford Handbook on Reasons and Normativity*.
- Helm, B.W. (2001). *Emotional Reason*. Cambridge: Cambridge University Press.
- Helm, B.W. (2009). Emotions as Evaluative Feelings. *Emotion Review*, 1(3), 248--55.
- Hornsby, J. (2007). Knowledge, belief, and reasons for acting. In C. Penco, M. Beaney, & M. Vignolo (Eds.), *Explaining the Mental* (pp. 88–105). Newcastle: Cambridge Scholars Publishing.
- Hornsby, J. (2008). A disjunctive conception of acting for reasons. In A. Haddock & F. MacPherson (Eds.), *Disjunctivism: Perception, Action, Knowledge* (pp. 244–261). Oxford: Oxford University Press.
- Hughes, N. (2014). Is knowledge the ability to act for the reason that P? *Episteme*, 11(4), 457–462.
- Hyman, J. (1999). How knowledge works. *Philosophical Quarterly*, 49(197), 433–451.
- Hyman, J. (2011). Acting for reasons: Reply to Dancy. *Frontiers of Philosophy in China*, 6(3), 358–368.
- Hyman, J. (2015). *Action, Knowledge, and Will*. Oxford: Oxford University Press.
- James, W. (1884). What is an emotion? *Mind*, 9 (34): 188–205.
- Kensinger, E. A., Garoff-Eaton, R. J., & Schacter, D. L. (2006). Memory for specific visual details can be enhanced by negative arousing content. *Journal of Memory and Language*, 54, 99 –112.
- Kensinger, E. A. (2007). Negative emotion enhances memory accuracy behavioral and neuroimaging evidence. *Current Directions in Psychological Science*, 16, 213–218.
- Kensinger, E. A. (2009). Remembering the details: Effects of emotion. *Emotion Review*, 1, 99 –113.
- Kramer, T. H., Buckhout, R., & Eugenio, P. (1990). Weapon focus, arousal, and eyewitness memory: Attention must be paid. *Law and Human Behavior*, 14(2), 167-184.
- Krusemark EA, Li W (2011) Do all threats work the same way? Divergent effects of fear and disgust on sensory perception and attention. *The Journal of Neuroscience* 31: 3429–3434.  
pmid:21368054
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
- Loewenstein, G., & Lerner, J. S. (2003). The Role of Affect in Decision Making. In R. J. Davidson et al. (Eds.), *Handbook of Affective Science* (pp. 619-642). Oxford: Oxford University Press.

- Lange, C. G. (1885). *Om sindsbevaegelser: et psyko-fysiologisk studie*. Kjbenhavn: Jacob Lunds.  
Reprinted in *The Emotions* (C. G. Lange & W. James, eds., I.A. Haupt, Trans.). Baltimore: Williams & Wilkins Company 1922.
- LeDoux, J. (1989). Cognitive-Emotional Interactions in the Brain. . *Cognition and Emotion*, 3 (4):267-289. .
- LeDoux, J. (1996). *The Emotional Brain*. New York: Simon & Schuster.
- LeDoux, J. (2000). Cognitive-emotional interactions: Listen to the brain. In L. N. Richard D. R. Lane (Ed.), *Cognitive Neuroscience of Emotion* (pp. 129--155). Oxford University Press.
- Lerner, J. S., & Keltner, D. (2000). Beyond valence: Toward a model of emotion-specific influences on judgment and choice. *Cognition and Emotion*, 14, 473-493.
- Lerner, J., Li, Y. and Valdesolo, P., Kassam, K.S., (2015) Emotion and Decision Making. *Annual Review of Psychology*, Vol. 66, pp. 799-823.
- Liu, Y., Zhang, D., & Luo, Y. (2015). How disgust facilitates avoidance: an ERP study on attention modulation by threats. *Social cognitive and affective neuroscience*, 10(4), 598-604.
- Locke, D. (2015). Knowledge, explanation and motivating reasons. *American Philosophical Quarterly*, 52(3), 215–232.
- Marks, J. (1982). A Theory of Emotion. *Philosophical Studies*, 42: 227–242.
- Maher, P. (2005). Confirmation Theory, *The Encyclopedia of Philosophy*, 2nd edition, Donald M. Borchert (ed.), Detroit: Macmillan.
- Morris, J.S., Friston, K.J., Buchel, C, Frith, C.D., Young, A.W., Calder A.J. Dolan, R.J. (1998). A neuromodulatory role for the human amygdala in processing emotional facial expressions. *Brain*, 121, 47–57.
- McDowell, J. (2013). Acting in the light of a fact. In D. Bakhurst, B. Hooker, & M. O. Little (Eds.), *Thinking about reasons: Themes from the philosophy of Jonathan Dancy* (pp. 13–28). Oxford: Oxford University Press.
- Nussbaum, M. (2001). *Upheavals of Thought: The Intelligence of Emotions*. Cambridge: Cambridge University Press.
- Nussbaum, M. (2004). Emotions as Judgments of Value and Importance. In R. Solomon (Ed.), *Series in affective science. Thinking about feeling: Contemporary philosophers on emotions* (pp. 183-199). NY: Oxford University Press.
- Partee, B. (1974). Opacity and Scope. In M. M. Unger (Ed.), *Semantics and Philosophy* (pp. 81 101). New York: NYU Press.
- Perry, J. (1979). The Problem of the Essential Indexical. *Noûs*, 13(1), 3-21. .
- Prinz, J. (2004). *Gut Reactions: A Perceptual Theory of Emotion*. Oxford: Oxford University Press.

- Pritchard, Duncan (2005). *Epistemic Luck*. Oxford University Press UK.
- Rabinowitz, D. (2011). *The Safety Condition for Knowledge* . Retrieved from Internet Encyclopedia of Philosophy: <https://www.iep.utm.edu/safety-c/>
- Ramsey, F. (1929/1990). General Propositions and Causality. In H. Mellor (Ed.), *Philosophical Papers*. Cambridge University Press.
- Richter-Levin, G., & Akirav, I. (2003). Emotional tagging of memory formation—in the search for neural mechanisms. *Brain Research Reviews*, 43, 247-256.
- Romanski L.M., L. J. (1992). Equipotentiality of thalamo-amygdala and thalamocortico-amygdala circuits in auditory fear conditioning. *Journal of Neuroscience* , 12:4501–9 .
- Roberts, R. C. (2003). *Emotions:an Essay in Aid of Moral Psychology*. Cabridge: Cambridge University Press.
- Robinson, J. (2004). Emotion: Biological Fact or Social Construction? In R. C. Solomon (Ed.), *Thinking About Feeling: Contemporary Philosophers on Emotions* (pp. 28–43. ). Oxford: Oxford University Press.
- Russell, B. (1905). On Denoting. *Mind*, 14(56), 479-493.
- Rozin, P., Fallon, A. E. (1987) A perspective on disgust. *Psychological Review*, 94, 32-41.
- Scarantino, A. (2010). Insights and Blindspots of the Cognitivist Theory of Emotions. *British Journal for the Philosophy of Science* 6 (4):729-768.
- Schwarz, N. (1990). Feelings as information: Informational and motivational functions of affective states. In E. H. Sorrentino (Ed.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 527-561). New York: Guilford Press.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61(2), 195-202.
- Schwarz, N. (2011). Feelings-as-information theory. In *Handbook of theories of social psychology* (Vol. 1, pp. 289-308).
- Solomon, R.C. (1977). *The Passions: The Myth and Nature of Human Emotion*. Anchor Press.
- Solomon, R.C. (1988). On Emotions as Judgments *American Philosophical Quarterly*, 25(2), 183-191.
- Solomon, R. C. (2003). *Not Passion's Slave*. Oxford: Oxford University Press.
- Sorabji, R. (2000). *Emotion and Peace of Mind*. Oxford: Oxford University Press.
- Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." In James Tomberlin (ed.), *Philosophical Perspectives* 13: Epistemology. Oxford: Blackwell, 141-54.
- Stalnaker, R. (1984). *Inquiry*. Cambridge: Cambridge University Press.
- Stalnaker, R. (1990). Mental Content and Linguistic Form. *Philosophical Studies*, 58(1/2), 129-146.

- Stanley, J. & Williamson, T. (2001). Knowing How. *The Journal of Philosophy*, 98(8), 411-444.
- Stanley, D. A., Ferneyhough, E., & Phelps, E. (2009). Neural perspectives on emotion: Impact on perception, attention and memory. In G. G. Berntson, & J. T. Cacioppo (Eds.), *Handbook of neuroscience for the behavioral sciences* New York: John Wiley and Sons.
- Teroni, F. (2007). Emotions and Formal Objects. *Dialectica*, 61(3), 395-415.
- Tiedens, L. Z., & Linton, S. (2001). Judgment under emotional certainty and uncertainty: The effects of specific emotions on information processing. *Journal of Personality and Social Psychology*, 81(6), 973-988.
- Todd, R., Cunningham, W., Anderson, A., & Thompson, E. (2012). Affect biased attention as emotion regulation. *Trends in Cognitive Sciences*, 16(7), 365–372.
- Turri, J. (2010). Does perceiving entail knowing? *Theoria*, 76(3), 197–206.
- Unger, P. K. (1975). *Ignorance: A case for Scepticism*. Oxford: Clarendon Press.
- Vuilleumier, P. & Schwarz, S. (2001). Beware and be aware: Capture of spatial attention by fear-related stimuli in neglect. *Neuroreport*, 2 (6), 1119-1122.
- Walton, K. (1978). Fearing Fictions. *The Journal of Philosophy*, 75 , 5-27.
- Whalen, P.J., Rauch, S.L., Etcoff, N.L., McInerney S.C., Lee, M.B., Jenike, M.A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience*, 18, 411–418.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2016a). Knowing and Imagining, In Kind, A., & Kung, P. (Eds.), *Knowledge Through Imagination*, Oxford: Oxford University Press.
- Williamson, T. (2016b). Counterpossibles in metaphysics. In Armour-Garb, B. & Kroon, F. (Eds.), *Philosophical fictionalism. Forthcoming*.
- Williamson, T. (2017). Acting on knowledge. In J. A. Carter, E. Gordon, B. Jarvis (Eds.), *Knowledge-first*, Oxford: Oxford University Press.
- Zhao, Z., Young, A., Jam, C., Fu, X., Zhang, W., Zhou, Q., Chen, W., Ni, X., Lin, J.K., Yang, J., Xu, X., Shen, W.L. (2017). A hypothalamic circuit that controls body temperature. *Proceedings of the National Academy of Sciences*. 114. 201616255.
- Zajonc, R. B. (1994). An early insight into the affect-perception interface. In P. K. Niedenthal (Ed.), *The Heart's Eye: Emotional Influences in Perception and Attention*. Academic Press.
- Zimmerman, T. E. (1993). On the Proper Treatment of Opacity in Certain Verbs. *Natural Language Semantics*, 1, 149–179.

