

Emotion, Epistemic Assessability, and Double Intentionality

Tricia Magalotti and Uriah Kriegel

Forthcoming in *Topoi* (special issue on double intentionality)

Emotions seem to be epistemically assessable: fear of an onrushing truck is epistemically justified, *mutatis mutandis*, whereas fear of a peanut rolling on the floor is not. But there is a difficulty in understanding *why* emotions are epistemically assessable. It is clear why *beliefs*, for instance, are epistemically assessable: epistemic assessability is, arguably, assessability with respect to likely truth, and belief is by its nature concerned with truth; truth is, we might say, belief's "formal object." Emotions, however, have formal objects different from truth: the formal object of fear is danger, the formal object of indignation is injustice, the formal object of grief is loss, and so on. After considering how a number of different accounts of emotion might account for the epistemic assessability of emotion, we develop a novel account of the domain of the epistemically assessable, according to which any mental state which is constitutively evidence-responsive is epistemically assessable, regardless of whether its formal object is truth.

1. The Epistemic Assessability of Emotions

Suppose you sit in a dark bunker during an enemy air raid and suddenly wonder whether outside the bunker the sun is shining and the birds are singing. Suppose further that you then decide to believe that this is so. (Assume for the sake of argument that this is something you can do). This belief may be *practically* justified insofar as it could keep you from falling into despair. Still, if there is no evidence that the sun is in fact shining and birds are in fact singing, your belief is not *epistemically* justified. For there is no information at your disposal to suggest the likely truth of the proposition 'The sun is shining and the birds are singing.'

Belief is thus *epistemically assessable*: it makes sense to ask, of any particular belief, whether it is epistemically justified or not. What does it take *in general* for a mental state to be epistemically assessable? Our vignette suggests the relevance of evidence possession and likely truth, but even without a general account of epistemic assessability, however, we can appreciate that some mental states are paradigmatically epistemically assessable while others are paradigmatically non-assessable. Belief is a paradigm of epistemic assessability: it is never a category mistake to say that some belief is epistemically justified or unjustified. Traditionally, perceptual experiences have been taken to be epistemically *non-assessable*, though the view has become more contentious of late (Siegel 2017). Still, *bodily sensations* certainly offer a paradigm of epistemic non-assessability. If you feel nauseous, or ticklish, it is hard to see in what sense your nausea/tickle experience could be epistemically justified or unjustified.

Emotions seem in this respect more like beliefs than like nausea: they can be legitimately assessed epistemically. In short:

Epistemic Assessability: Emotions are epistemically assessable.

We formulate Epistemic Assessability as a generic rather than as a quantified proposition: it is not essential to our purposes here to assert the *universally* quantified proposition that all emotions are epistemically assessable, though we do not wish to make a simple existential claim that might be supported by a handful of atypical cases. What Epistemic Assessability asserts is that in many perfectly normal, in no way atypical cases, emotions are epistemically assessable.

We will offer a more complete defense of Epistemic Assessability in section V. For now, we just want to note that Epistemic Assessability seems to be presupposed by our social practices surrounding emotion.

First of all, we clearly have a social practice of *rationally criticizing* one another's emotional states (Salmela 2006), and the rationality at issue is often epistemic. People often accuse one another of having irrational emotions, such as anger at a small child for accidentally spilling a glass of milk or disgust toward people of a particular race. In some cases, the sort of irrationality at issue may be practical, but in others it is surely epistemic. Compare "her jealousy is so irrational: in the end, it will only tear her and her partner apart" and "his jealousy is so irrational: all his partner did was smile at the waiter." The first critique is practical, but the second is epistemic. Clearly, we take emotions, at least sometimes, to be subject to norms of epistemic rationality.

Conversely, we also have a practice of offering (epistemic) reasons for our own emotions, and we tend to withdraw our accusations of irrationality if someone gives us sufficiently good reasons. Suppose Alice, Bob, and Cat are talking about baseball. Cat says, "Did you hear that the Red Sox won the game yesterday?" Bob responds, "I know, I'm so happy!" Alice says "Why are you happy about that? You don't even like baseball." Alice is challenging Bob's emotional response. She does not think that it makes sense for him to be happy that some particular team won a game given that he does not care about baseball. But suppose Bob replies "You're right, I don't like sports. But my cousin plays for the Red Sox, and that is why I am happy." At this point, Alice would concede that it is reasonable for Bob to be happy about the win. Importantly, when Bob says that his cousin plays for the Red Sox, he offers a specifically epistemic reason for his happiness that the Red Sox won: he provides evidence supporting the claim that the win is good for him.

If we denied that emotions were epistemically assessable, the practices of (epistemic-)rational critique and (epistemic-)reason giving surrounding emotion would be misguided. But they do not seem misguided at all. In this respect, emotions seem more like beliefs than bodily sensations. It is hard to imagine a community of creatures who criticize each other—epistemically—for feeling nauseous or hungry or crampy, and who then defend themselves by giving reasons—epistemic reasons—to be so. *That* community's practice seems misguided in a way ours does not.

2. The Double Intentionality of Emotion

According to a popular model in the philosophy of emotions, emotions have in some sense a *double intentionality*: each emotion is intentionally directed toward two different kinds of object. One is the “particular object”: the thing that, colloquially, we would say the emotion is *about*. When one is afraid of a snake, the snake is the particular object of one’s fear. Particular objects can be concrete individual objects, such as snakes, but they may also be propositional objects, as in the fear *that* the snake will hurt me. In addition, emotions are also thought to have a second kind of object, which has come to be known as their “formal object” (Kenny 1963). These are thought to be properties at which entire emotion *types* are in some sense aimed. Fear is aimed at dangerousness, grief at loss, and so on.

One way to think about the relationship between particular and formal objects is in terms of a certain division of individuating labor. If we ask of a token fear experience *F*, (a) what makes *F* the fear it is and (b) what makes it a fear at all, the answer is (i) that *F*’s *particular* object is what makes it the fear it is and (ii) its *formal* object is what makes it a fear at all. Joe’s fear of the snake is the fear it is because the snake is its particular object, and it is a fear at all because dangerousness is its formal object. In other words, particular objects are responsible for the individuation of *token* emotions, while formal objects are responsible for the individuation of emotion *types*. It is in this sense that Joe’s token fear is aimed at a snake whereas fear as such is aimed at dangerousness. The double aim, or double intentionality, of emotions reflects their double metaphysical status as token experiences and as tokening experience types.

There is also a division of labor here with respect to *correctness conditions*. Emotions are sometimes correct and sometimes incorrect. Joe’s fear of the snake is correct (in the sense that Joe’s fear gets things right) just if the snake is dangerous. Here the particular object, the snake, tells us what we need to evaluate, while the formal object, dangerousness, tells us what kind of evaluation we must perform, to establish whether Joe’s fear is correct. Evaluating fears for correctness is just evaluating feared objects for dangerousness. In general, an emotion’s formal object determines the kind of evaluation relevant to whether the emotion is correct, while its particular object provides the item on which the evaluation is to be made.

What exactly is the relationship between formal objects and particular objects? Are they both intentional objects? If they are, what distinguishes the formal from the particular? If they are not, what role does each play within the overall intentional structure of a mental state?

One epistemic possibility is that formal objects are indeed intentional objects, in the same sense particular objects are, but that they are intentional objects common and peculiar to certain mental state types. On this view, when we say ‘the formal object of fear is the dangerous,’ we mean roughly the following:

(FO1) The property of being dangerous is represented by all and only fears.

However, it seems false that *only* fears represent dangerousness. A person may *believe* that, say, scuba diving is dangerous—regardless of whether she is afraid of it. And she may also *wonder* whether scuba diving is dangerous, *hope* that it is dangerous (for whatever reason), and so on.

A second possibility is that formal objects are the *explanatory grounds* of the subject’s experiencing an emotion. On this view, when *S* is afraid of a dog, the representational content of *S*’s fear is the dog, not the dog’s being dangerous; nonetheless, it is *because* the dog is dangerous

(or taken by S to be dangerous) that S is in a state of fear. Thus the dog's being dangerous (or its being *taken* by S to be dangerous) is the *reason why* S is in a state of fear. Here 'the formal object of fear is the dangerous' means something like:

(FO2) The property of being dangerous is the explanatory ground for all and only fears.

The main disadvantage of FO2 is that it casts formal objects as external to mental states' intentional structure. If danger is just the explanatory ground of fear's representation of the dog, it is not part of fear's representational character itself, but something that lies outside it. There is, on this view, no double intentionality strictly speaking. There is just intentional directedness at a dog, and a reason *why* the dog becomes the intentional object of a fear. The problem with this is that dangerousness does seem to characterize the *intentionality* of fear. Fear *frames* the world a certain way, namely, as involving dangers. The world of the fearful is unlike the world of the fearless.

A third suggestion is that formal objects are not what *explains*, but what *justifies*, or *rationalizes*, the relevant mental states. It is *fitting, rational, appropriate* to fear that which is dangerous. Accordingly, the correctness conditions of fear advert to the dangerous. Here 'the formal object of fear is the dangerous' means:

(FO3) The property of being dangerous provides the correctness conditions for all and only fears.

The main problem for FO3 is with explanatory depth. FO3 offers a *normative* gloss on formal objects. But normative characteristics supervene upon non-normative characteristics, which (we tend to think) underlie them and are explanatorily prior to them. Thus, we should expect fear to have some non-normative, *psychological* characteristics that underlie the fact that it is fitting to fear only the dangerous. It would be odd to treat as brute and inexplicable the fact that fears are fitting only when directed at the dangerous. More plausibly, there is something about what fear is *like* that *makes* it fitting to experience fear toward the dangerous.

Our own suggestion is that talk of formal objects is in fact a roundabout way of talking about the *mode* or *attitude* of mental states. Each emotion type employs a distinctive mode of intentionally relating to its intentional object, a distinctive attitude it takes toward its content (what it represents). As we understand these attitudinal features, they are dimensions of mental states' intentional structure, and are in fact ways of *framing* their intentional objects. They are just not part of that-which-is-being-represented. Thus, fear of a dog does not represent the dog as dangerous, insofar as danger is not part of what the fear represents; but it does represent-*as-dangerous* the dog, insofar as danger characterizes the manner in which the dog is framed in the fear. On this view, 'the formal object of fear is the dangerous' means:

(FO4) All and only fears involve an attitudinal feature of representing-as-dangerous.

By the same token, grief represents-as-a-loss its object, indignation represents-as-unjust its object, and so on. We contend that FO4 has many of the advantages, and none of the disadvantages, of FO1-FO3. First, FO4 has FO1's advantage of casting formal objects as a dimension of intentionality, but like FO2, it does not construe formal objects as part of what a mental state represents, and like FO3, it seems to get the extension right: it is hard to imagine a case where someone has a state that represents-as-dangerous something the person does not fear, or where the person fears something none of their mental states represents-as-dangerous. At the same time,

unlike FO1, FO4 designates as the formal object of fear something that no other mental state features; unlike FO2, FO4 makes formal objects part of mental states' intentionality, rather than something external to intentionality that merely *explains* of the state's intentionality; and unlike FO3, FO4 identifies formal objects with psychological features rather than with brute normative properties.

For all these reasons, we prefer (FO4). From this perspective, talk of double intentionality comes down to highlighting the Janus-faced character of intentionality as involving both intentional objects or contents (what is represented by those states) and intentional mode or attitude (how what is represented is framed in those states). As noted, however, this particular take is not strictly mandatory for the arguments to follow. We include this discussion mostly to give more texture to some of the key concepts we will be using.

3. Formal Objects and the Scope of the Epistemic

The discussion of formal objects in the emotion literature is reminiscent of the notion that mental state types have an *aim*. It is commonly held, for instance, that *the aim of belief is truth*.¹ This is naturally understood as the claim that the formal object of belief is truth. Thus, Petra's belief that the weather is nice is the belief it is because of the particular object it is aimed at, the proposition that the weather is nice, but is a belief at all because of the formal object it is aimed at, truth. And her belief is correct (she is right to believe this) just if the proposition that the weather is nice is true. Truth is what we need to evaluate the proposition *for* when we evaluate Petra's belief for correctness.

There are two different ways, however, of understanding correctness talk (see Glüer and Wikforss 2009), and correspondingly two different ways of understanding aim talk. On a purely *descriptive* understanding, correctness and incorrectness amount to something like accuracy and inaccuracy. On a more *normative* understanding, to say that a belief is correct is not merely to describe it, but also to commend it. The idea is that true beliefs are *successful* whereas false ones are *defective*. A true belief is a *good instance* of belief, an untrue belief is a *bad instance*.

By the same token, on a descriptive understanding, to say that truth is the aim of belief is to say something about the *nature* of belief: perhaps precisely that its formal object is truth. On a normative understanding, to say that truth is the aim of belief is to say something about the *norm* of belief: perhaps that we should manage our beliefs with an eye to promoting truth among them.

Now, it is natural to think that the *norm* of belief is grounded in the *nature* of belief: if belief were a different kind of thing, it would be governed by different norms. It is because truth is the formal object of belief that we ought to promote truth among our beliefs. There are, of course, many ways in which truth may be promoted, but the norm that we should pursue these ways will always be grounded in the fact that truth is the formal object of belief. Take, for instance, the basic maxim *believe that p only if your evidence supports p*—or as Clifford (1877, 295) memorably put it, “it is wrong always, everywhere, and for everyone, to believe anything upon insufficient evidence.”

¹ This is not the only possible view in this area, of course. Alternatives include that belief is aimed at knowledge (Williamson 2000) or justification (Feldman 2005). But here we will assume that truth is the aim of belief.

Arguably, it is because truth is the aim of belief in the descriptive sense (i.e., is the formal object of belief) that beliefs are subject to this maxim. More generally:

Aim-Norm Connection: It is in virtue of the fact that truth is the formal object of belief that beliefs are subject to truth-promoting epistemic norms.

Aim-Norm Connection is widely accepted in the aim-of-belief literature (see, e.g., Velleman 2000, 246 and Boghossian 2003, 39).

Aim-Norm Connection is often motivated by appeal to the notion of *regulation*. According to Velleman (2000), for instance, the fact that beliefs are regulated for truth grounds both the fact that truth is the standard for normative correctness of belief and the fact that beliefs are subject to evidential norms. This regulation can be understood in at least two ways: (a) in a personal-level, agential sense, whereby the epistemic agent attempts to believe a proposition “with the aim of thereby accepting a truth” (ibid., 251); or (b) in a sub-personal, “procedural” sense, whereby the aim of accepting a truth is implicit in the cognitive system’s workings and encoded in its general tendency to form beliefs that are somehow indicated to be true.

Now, if Aim-Norm Connection is true, it may seem but a short step to the claim that beliefs have a *unique* epistemic status in virtue of having the formal object of truth. If what grounds the fact that belief is subject to epistemic norms is the fact that truth is the formal object of belief, then it is natural to expect mental states the formal objects of which are *not* truth to be subject to *different* norms, or at any rate to not be subject to epistemic norms (and thus not be epistemically assessable). Thus we obtain the following principle:

Epistemic Exclusivity: Only mental states whose formal object is truth are epistemically assessable.

Because beliefs have truth as their formal object, they aim at truth and are regulated for truth. Their aiming at, and being regulated for, truth makes it the case that they are subject to epistemic norms requiring that we believe in truth-conducive ways – that they are, in short, epistemically assessable. Mental states that have a different formal object, now, will *not* be regulated for truth but for something else. They will not aim at the truth but at whatever their formal object is.

Another motivation for Epistemic Exclusivity has to do with the foundational role of truth in epistemology. According to “veritists,” true belief is the only thing of fundamental epistemic value. It is not uncommon to think that epistemic norms are grounded in or explained by this value of true belief. For example, one might think that it is epistemically required that we believe in accordance with the evidence *because* doing so either promotes or respects the epistemic value of true belief. If veritism is true, and if epistemic norms stand in this special relationship to fundamental epistemic value, then one might be led to think that beliefs are the only mental states that are subject to these truth-grounded epistemic norms. Other mental states, with other formal objects, will fall outside the scope of epistemic normativity.

Epistemic Exclusivity spells trouble, however, for the principle of Epistemic Assessability (i.e., emotions are epistemically assessable). For as we have seen in section II, the formal objects of emotions are *not* truth. They are dangerousness (for fear), injustice (for indignation), loss (for grief), and so on – but not *truth*. But if Epistemic Exclusivity is true, then emotions do have truth as their formal object or emotions are not epistemically assessable.

In other words, what we have on our hands here is an *inconsistent triad*. At least one of the following three antecedently plausible theses must be rejected:

Epistemic Assessability: Emotions are epistemically assessable.

Epistemic Exclusivity: Only mental states whose formal object is truth are epistemically assessable.

Not Truth: The formal object of emotions is not truth.

The question is which one to reject. In section IV, we consider the option of rejecting Not Truth, and in section V, that of rejecting Epistemic Assessability; we argue against both. Then, in section VI, we present our own approach to the problem, which consists rather in rejecting Epistemic Exclusivity.

4. Rejecting *Not Truth*: Cognitivist Approaches

The most natural way to reject *Not Truth* is to adopt a broadly cognitivist or judgmentalist account of emotion. According to “judgmentalists” (Solomon 1976, Nussbaum 2001), emotions are in fact beliefs, or judgments – doxastic states that are truth-apt. Fear of a snake, for example, just is the belief, or judgment, that the snake is dangerous, while grief over the death of a grandparent just is the belief that the grandparent’s death is a loss (or a significant personal loss, or whatever).

On this approach, emotions are not a *sui generis* kind of mental state but a particularly important subset of beliefs (important enough, say, for folk psychology to produce specialized labels for them). What distinguishes fears from griefs, on this view, is just the predicative component of the relevant belief contents: fear-beliefs predicate dangerousness of something, while grief-beliefs predicate loss-ness of something. Crucially, then, the difference between fear and grief pertains not to their formal but to their *particular*, propositional object. The formal object of all emotions is the same: truth. If so, Not Truth can be rejected.

The problem, however, is that the judgmentalist account of emotion strains credulity. At a very intuitive phenomenological level, emotional experience seems to be “affectively hot” in a way belief is not.² But more deeply, coming down on an issue and making a judgment *feels* like committing to the truth of some proposition. Consider this example from Eli Chudnoff (2015: 98): is it true that if $a < 1$, then $2 - 2a > 0$? If you are like us, you take a few seconds to think through this and then the feeling of truth descends on you. In this example, the feeling of truth-acceptance comes into phenomenological relief because of the element of mental effort required and the slight time-delay in experiencing it. Arguably, however, the feeling is there, albeit in a more subdued fashion, in all cases of consciously judging that something is the case. This feeling, however, clearly does not exhaust the phenomenology of typical *emotional* experiences. In grief over a loved one’s death, there may be a feeling of truth-acceptance toward the proposition ‘Grandma died.’ But

² Nussbaum (2004) argues that the “hotness” of emotions is due to the fact that the relevant judgments concern, or are about, matter of special personal importance to the person doing the judging. We find this somewhat underwhelming: applying the concept of personal importance may *reflect* the hotness of emotion, but does not seem to *constitute* this hotness.

this does not yet capture the loss-y feeling of the grief, and that loss-y feeling is not simply a feeling of truth-acceptance where the proposition accepted as true is ‘Grandma’s death is a terrible loss.’ There is more to the loss-feeling characteristic of grief than a feeling of truth-acceptance. Similarly, when you suddenly notice an onrushing truck coming from your blind side, and consequently experience an upheaval of fear, what you experience feels very different from the upheaval of truth-recognition or truth-acceptance that comes through in Chudnoff’s case.³

Some approaches to emotion are not quite as radical as to *identify* emotions with doxastic states, but still claim that emotions *include* doxastic states as constituents. We call these *doxastic-constituent views* (DCVs), and use the label “cognitivism” to cover both judgmentalist and DCVs. DCVs are also naturally seen as rejecting Not Truth, since they claim that an emotion is at least in part a mental state that has truth for its formal object.⁴ We will now consider three different DCVs, arguing against each.

Perhaps best known among DCVs is the view that emotions are belief-desire compounds (Marks 1982, Gordon 1987). On this view, fear of a snake is a combination of a belief that the snake is dangerous and a desire to flee the scene. When these co-occur in you, you experience fear, not because they *cause* fear in you, but because that is just what experiencing fear *is*.

There are various problems with the belief/desire compound view, of which the most important for present purposes is this. Consider again the experiential element in fear of an onrushing truck that seemed incongruent with a phenomenological description in terms of accepting the truth of a proposition. That additional element characteristic of fear does not seem to be what is recovered when we add the phenomenology of desire to the phenomenology of belief or judgment. It is of course difficult to pin down what the distinctive phenomenology of desire exactly is, and we cannot hope to consider all theoretical options here. But with the fear of the onrushing truck there seems to be some kind of vivid feeling of being pulled into action to avoid a negative result: you feel vividly pulled to jump out of the way. But adding this feeling to the feeling of accepting the truth of the proposition that the truck is dangerous does not quite seem to make up the difference with the complete phenomenology of fear. There is an affective quality to fear, pegged precisely to the felt sense of danger or looming harm, that goes beyond that felt pull to action. Indeed, as you notice the onrushing truck, you may find yourself *paralyzed* by fear rather than moved to action; your fear as the truck closes in on you need not subside in consequence, however.

A second DCV would add to belief not desire but something like the proprioceptive or kinesthetic awareness, or feeling, of bodily changes, often in internal organs. In Schachter and Singer’s (1962) theory, for instance, an emotion is construed as a combination of two factors, one cognitive and the other a felt physiological arousal. Fear, on this view, consists in a feeling of a

³ We are open to the possibility of atypical cases of emotional experience, perhaps of the sort designated sometimes as “metacognitive feelings,” that do incorporate a phenomenal element of truth-acceptance. But this would not solve the problem as applied to the more familiar emotions mentioned in the main text.

⁴ More precisely, there are two ways this approach could play out as a response to the problem posed by the above inconsistent triad. In one version, the claim would be that emotions have two formal objects: truth and whatever the formal objects are of the other mental-state constituents of emotion. In this version, Not Truth is certainly denied. In a different version, the claim would be that emotion has a single sui generis formal object that has truth as a constituent. In this version, the view would deny Epistemic Exclusivity, claiming that to be epistemically assessable a mental state need not have truth for its formal object, as long as it has truth for a *constituent* of its formal object. Either way, the problem would be solved and Epistemic Assessability would be saved.

pounding heart etc. *doxastically interpreted* as due to an impending danger. A third DCV adds to belief *both* desire and bodily feeling (Kriegel 2014). The advantage of adding bodily feelings to the account, whether or not one preserves desire, is that bodily feelings seem much more apt for capturing the affectively hot element of the fear experience. The sensations of visceral tumult, racing heart, and so on are at least the kind of thing such that, once we imagine them added to the phenomenology of accepting-as-true, we get a sense of the properly affective upheaval characteristic of experiencing fear. At the same time, the view preserves a doxastic constituent in emotion that might explain its epistemic assessability.

This is clearly DCV's most plausible version. We do worry that construing fear as involving the belief that something is dangerous, whatever else one adds into the mix, would require the fearing subject to possess the concept of dangerousness. It is empirically quite plausible, it seems to us, that the capacity for fear predates the possession of this concept in infants. As Morreall (1993: 361) points out, "Infants fear loud noises long before they have . . . concepts like danger." And similarly for other emotions: a toddler's capacity for indignation predates her possession of the concept of injustice, her capacity for grief predates her possession of the concept of loss, and so on.

However, our *main* critique of the view under consideration targets specifically the *epistemology* of emotion. We want to argue that the view faces two important difficulties, difficulties that go to the very idea of reducing the epistemic assessability of emotions to that of beliefs.

First difficulty. Although there may be something epistemically problematic about holding inconsistent emotions, it is not quite the epistemic disgrace that holding inconsistent beliefs represents. Suppose you apply for a grant and the only other applicant—this you know—is your best friend and much admired colleague. She gets the grant. You are happy that she got the grant and you are sad that you did not – even though her getting the grant and your not getting it are one and the same state of affairs, and, crucially, you *know* that they are (Montague 2009). If being happy about something is just believing that it is good for you, and being sad about something is just believing that it is bad for you, then here you hold inconsistent beliefs: that the relevant state of affairs is both good for you and bad for you. It is true, of course, that the state of affairs is represented in these two beliefs under different modes of presentation. But you are *aware* that the two modes of presentation are co-referential. This is like believing that Phosphorus is round and that Hesperus is square *despite knowing that Phosphorus is Hesperus*. It is an inconsistency pure and simple. Now, although being both happy and sad about one and the same state of affairs is not epistemically irreproachable, it is more of an epistemic misdemeanor than an epistemic felony; whereas holding inconsistent beliefs is a cardinal epistemic sin. If emotions were epistemically assessable purely in virtue of having belief constituents, such a gap in epistemic status would be incomprehensible.

Might the cognitivist bank on beliefs sufficiently intricate that holding them together would be more like the epistemic misdemeanor we have in mind? One option here is to rely on *contrastive beliefs* as the doxastic constituents of emotion. For instance, being happy that your best friend got the grant might be taken to involve as constituent the belief "it is good for me that my best friend rather than some stranger got the grant," while being sad that you did not get the grant involves as constituent the belief "it is bad for me that my best friend rather than I got the grant." There might be some tension in holding seriously both beliefs, but clearly they are not strictly inconsistent.

In response, we note that there is something psychologically unrealistic in requiring emotions to involve such intricate beliefs as constituents. On the view under consideration, a person who simply believes that x is bad for them is not emoting yet, but once they think that x *rather than* y is bad for them, they are emoting. Sadness separates itself from mere belief partly in virtue of the contrastive structure of its content. We find this psychologically unrealistic as an account of the nature of emotion. Moreover, we suspect again that infants and animals may be able to emote long before they can entertain contrastive or other specially intricate propositions.

Second difficulty. On the face of it, the *strength* of an emotion is also epistemically assessable. If an unleashed midsized dog approaches you, mild apprehension may be epistemically justified, but frenzied terror would be epistemically unjustified. And when we fear something which is not at all dangerous, the more intense our fear the *more epistemically unjustified* it is (being petrified of a peanut is epistemically worse than being mildly apprehensive about it). Thus our emotional experiences are epistemically assessable not only for their qualitative but also for their *quantitative* properties. Now, if the epistemic assessability of emotions reduced to that of belief, this would have to be explained by changes in strength of the emotion's belief constituent. But as we will now argue, it is hard to make this sound remotely plausible.

One cognitivist idea might be that the affective intensity of an emotion is a matter of the degree of value predicated in the relevant belief. For one's fear of x to be more intense than one's fear of y is for one to believe that x is dangerous to degree $D+$ and believe that y is dangerous to degree $D-$, such that $D+ > D-$. This is problematic, however, insofar as the affective intensity of emotions seems to lie on a continuum, whereas belief contents are perforce discrete. That is, it is psychologically implausible, perhaps impossible, that we should possess a concept for each degree of dangerousness. To capture degree differences, our danger-beliefs must use modifier concepts such as "very" and "very very." But these will just have nowhere near the granularity needed to capture psychologically real differences in emotional intensity.

Another cognitivist idea might be that the relevant belief contents are the same, but are held with different degrees of confidence—and it is in virtue of this that they are differently evaluated. Thus, being very confident that the dog is dangerous is epistemically unjustified while suspecting that it might be dangerous is justified. This approach at least captures the continuous nature of emotional intensity. However, when we focus on the difference between mild apprehension and frenzied terror that explains their differential epistemic evaluation, what we focus the mind on is a difference in the affectively hot dimension of the emotion, not the cold doxastic dimension. It is the variation in *affective intensity* that we epistemically assess in the first instance: what is epistemically unjustified is being *this agitated* about the dog. It may well be, of course, that we can *also* epistemically assess doxastic strength (i.e., degree of confidence). But *in addition*, there is the epistemic assessment of the subject's level of affective agitation.

We take the moral of these two phenomena—that doxastic inconsistency is epistemically worse than emotional inconsistency and that affective intensity is epistemically assessable as well—to be that, regardless of whether emotions have doxastic constituents, emotions are epistemically assessable purely in virtue of their properly affective dimension. If this is right, then attacking Not Truth is the wrong approach to the puzzle captured in our triad.

5. Rejecting *Epistemic Assessability*: A Perceptualist Approach

What, then, about the option of rejecting Epistemic Assessability, the idea that emotions are epistemically assessable as justified or unjustified? Here the most natural approach is to adopt a broadly *perceptualist* view of emotion—given that most philosophers take perceptual experiences to be epistemically non-assessable.

It is not straightforward exactly what perceptualism about emotion *is*. On one extreme is the view that emotions *are* sense perceptions. More often, the claim is rather that emotions are *analogous* to perception in some important way. What that way is also varies from one version of perceptualism to another. Given the present paper's concern, the kinds of perceptualism that interest us are those which imply that emotions are analogous to perception *epistemologically*. Accordingly, we will use the label "perceptualism" for any view of emotion according to which emotions are like perceptions in all epistemic respects.

Defenders of perceptualism emphasize similarities between perceptual illusions and recalcitrant emotions. With both perceptual illusions and recalcitrant emotions, one's experience of the relevant object fails to match up with a judgment that one makes about the object. In the Müller-Lyer illusion, for instance, perceptual experience presents the lines as being different lengths despite one's judgment that they are the same length. Similarly, fear of flying can present flying as dangerous despite one's judgment that flying is safe. Perceptualists claim that just as visual illusions are not assessable as *epistemically unjustified*, neither are recalcitrant emotions. This point is then leveraged in support of the more general claim that *no* emotions are epistemically assessable, either as justified or as unjustified. If so, Epistemic Assessability is false.

To use perceptualism this way—i.e., to get out of the inconsistent triad by rejecting Epistemic Assessability—we need to assume that perception in general is not epistemically assessable. Since perceptualism claims that emotions are like perceptions in all epistemic respects, it would then follow that emotions are generally not epistemically assessable either. Now, as noted already in section I, it is possible, though quite unusual given such phenomena as the Müller-Lyer illusion, to hold that perceptual experience *is* epistemically assessable. If a perceptualist about emotion holds this view of perception, then the resulting version of perceptualism would *not* be the kind that could solve the inconsistent triad by rejecting Epistemic Assessability. Since our focus here is on solving the inconsistent triad from section III, we consider the more standard version of perceptualism about emotion, the one that assumes that perception in general is not epistemically assessable. If perceptualism of the sort we are interested in is true, then emotions are not epistemically assessable. In this way, perceptualists can maintain that the formal object of emotions is something other than truth (Not Truth) and that only mental states with the formal object of truth are epistemically assessable (Epistemic Exclusivity). Problem solved.

However, there are several problems with this solution to the problem. First, the perceptualist intuition in response to cases of recalcitrant emotion is not universally shared. To many, including us, it seems false that there is nothing epistemically objectionable about an emotional response that persists to eternity despite the fact that it conflicts with the subject's evidence. One standard line on these cases is that while simultaneously fearing flying and believing that flying is safe is *less* irrational than believing that flying is dangerous and believing that flying is safe, it is irrational nonetheless. In contrast, there is nothing irrational about simultaneously

believing that the Müller-Lyer lines are of equal length and having a perceptual experience as of their being the same length. There is thus an important epistemic disanalogy between perception and emotion (see Helm 2001 and following him many others).

Second, even if perceptualism did get things right about recalcitrant emotions, the emphasis on recalcitrant emotions obscures the commonsense observation that in many normal, in no way atypical cases, when we acquire evidence that undermines a particular emotional response, the emotional response tends to fade. That is, emotions are typically *not* recalcitrant. Suppose a man bumps into you on the subway. You feel inside your coat pocket and realize that your wallet is gone. Thinking that the man has stolen your wallet, you become angry. You then see that the wallet has actually fallen on the floor, and the anger dissipates rather quickly (even if not *immediately*). This is not some weird fringe case but a very typical sort of case: first, you acquire some evidence (the man bumped into you and the wallet is missing); second, an emotion is formed on the basis of that evidence (you become angry); third, you acquire some counter-evidence (the wallet is on the floor); finally, the emotion dissipates in light of the counter-evidence (you are no longer angry). Anger is thus often evidence-responsive in a way perception is not.

One might doubt whether in this case the emotion itself is responding to the evidence. Perhaps it is in the first instance your belief about the wallet being stolen that is responding to the evidence (e.g., the visual experience of the wallet on the floor), and the change in the emotion is merely a causal consequence of the revised belief. Maybe so. But this would still be a way for the emotion to respond to evidence. We might call it “doxastically mediated evidence-responsiveness.” Importantly, perceptions generally seem to lack even doxastically mediated evidence-responsiveness. Even when one confidently believes that the stick is straight, one still sees it as bent. Thus, the contrast between the evidence-responsiveness of emotions and perceptions remains.

A third problem with the proposed perceptualist resolution of our inconsistent triad is that the evidence-unresponsiveness of recalcitrant emotions may be a local failure rather than constitutive of emotion as such. Recall that the temptation to think that recalcitrant emotions are not epistemically assessable comes from their being unresponsive to evidence. This, after all, is what recalcitrant emotions have in common with visual illusions. But notice that local unresponsiveness to evidence does not establish that a state is epistemically non-assessable. Beliefs that do not respond appropriately to evidence, for example, may be epistemically unjustified (i.e., assessable negatively) rather than epistemically non-assessable. Astrological beliefs often fail to respond to evidence, but we do not exempt them from epistemic evaluation on that basis. On the contrary, we evaluate them specially harshly. In other words, there is a difference between local failure to respond to evidence and constitutive evidence-unresponsiveness. And as we have just seen with the wallet case, emotions do not seem to be *constitutively* evidence-unresponsive at all.

Recall, moreover, that Epistemic Assessability is not a *universally quantified* thesis. Thus for our purposes it is not essential that all emotions be epistemically assessable.⁵ As we have argued, in many very normal cases, when we acquire evidence that undermines a particular

⁵ We do recognize that there would be something theoretically unappealing about a picture of emotion in which some of them are epistemically assessable and some are not. So we face a choice between pushing for the universal thesis and explaining (away) a fairly fundamental disunity in emotion. This is an important choice, but we do not need to take a stand on it here. We suspect that ultimately all emotions are epistemically assessable, and the stubbornly recalcitrant ones are simply epistemically unjustified. But this position is not germane to our main concerns in this paper.

emotional response, that response tends to fade. Emotions may be *slower* than beliefs to respond to evidence, in that a certain emotional inertia often lingers with us beyond the moment of evidence acquisition. Nonetheless, emotions that are *eternally* and *absolutely* insensitive to evidence in the way nausea, hunger, and the perception of a bent stick are do not seem to be the rule; and all we want to insist on is that evidence-sensitive emotional responses are not fringe exceptions either.

This difference in evidence-sensitivity between perceptions and emotions points to a corresponding difference in their epistemic properties. The fact that perceptual experiences (like bodily sensations) are insensitive to changes in evidence is naturally taken to *explain* why they are not candidates for epistemic justification. Just as we take the actions of creatures irresponsible to moral reasons to be not morally evaluable, so too we should think that kinds of mental states that are irresponsible to epistemic reasons are not epistemically evaluable. Perceptions are not, in the typical case, responsive to reasons. But emotions are. And so, there is a principled difference between perceptions and emotions that warrants epistemically assessing the latter but not the former.

We conclude that the perceptualist interpretation of what is going on epistemologically with recalcitrant emotions is unconvincing.

6. Formal Objects and Epistemic Assessability

In this section, we offer our own solution to the problem posed by the inconsistent triad from section III, a solution consisting in the rejection of Epistemic Exclusivity. Our approach builds on three basic ideas. First, there is a distinction between the *what* and the *how* of regulation; in the case of belief, for instance, truth is *what* it is regulated for, and evidentially is at least part of *how* belief is regulated. Second, what makes belief the paradigm epistemically assessable state is not, at least not *just*, the fact that it is regulated for truth; it is also, at least in part, the fact that it is regulated evidentially. And third, although emotions are not regulated for truth, but for their own proprietary aims or formal objects, they are nonetheless *evidentially* regulated. *This*, we argue, is what makes emotions epistemically assessable: one can meaningfully ask of each token emotional experience whether it was formed in an evidentially proper way. And so Epistemic Exclusivity is to be rejected: Not only mental states whose formal object is truth are epistemically assessable. So long as a mental state type is *evidentially* regulated for whatever it is regulated for, epistemic evaluation of that state will make sense.

Start with the distinction between the what and the how of regulation. It is easy to appreciate *in abstracto* that the regulation of any system has a certain aim it tries to serve, and also employs a specific manner of attempting to best serve that aim. And in principle, two systems regulated for the same thing may employ different styles of regulation: one system regulates for aim A in way W_1 while another regulates for A in way W_2 . Take living organisms. These are complex systems (self-)regulated for something like survival and reproduction. But the natural world offers a staggering multiplicity of *ways* living organisms are (self-)regulated for survival and reproduction. In plants the system's activity is regulated centrally using photosynthesis; in animals this has no role. Conversely, it is also possible to imagine two complex systems regulated for different aims but exhibiting an overlap in the way they regulate for these aims. Thus, two systems of government, a liberal democracy and an autocratic dictatorship say, may have two completely different aims – one

is regulated for the wellbeing of the citizenry, whereas in the other that has no role – but the use of a police force in law enforcement is central to the *way* both systems regulate for their respective aims.

Naturally, the what/how distinction applies to the regulation of beliefs as well. In particular, we want to suggest, belief is not just regulated *for truth*, it is *evidentially* regulated for truth. That is, it is through the use of evidence that we regulate for truth. (It is an open question whether evidential regulation is the *only* form of belief regulation. We need not take a stand on this here.) Now, in the aim-of-belief literature, the notion of regulation-for-truth and the notion of evidential regulation tend to be run together, presumably due to insensitivity to the fact that the former concerns the *what*, or *aim*, of regulation, whereas the latter concerns the *how*, or *manner*, of regulation. Consider the following two passages from Shah and Velleman:

In forming and retaining a belief... one responds to evidence and reasoning in ways that are designed to be truth-conducive. Hence belief is regulated for truth, whereas other, non-belief-involving cognitive attitudes are not. (2005, p. 498)

Being regulated for truth consists partly in responsiveness to evidence. The belief that *p* tends to be formed in response to evidence of *p*'s truth, to be reinforced by additional evidence of it, and to be extinguished by evidence against it (*ibid.*, p. 500)

We agree with Shah and Velleman both that beliefs are regulated for truth and that they are evidentially regulated. But from this neither of the following follows:

(T1) If a mental state *M* is regulated for truth, then necessarily, *M* is evidentially regulated.

(T2) If a mental state *M* is evidentially regulated, then necessarily, *M* is regulated for truth.

We suspect that (T1) is nonetheless true. What we want to deny is (T2). We want to argue that emotions are evidentially regulated, despite not being regulated for truth. Thus there can be evidential regulation for aims other than truth.

Once we cleanly distinguish between regulation for truth and evidential regulation as two characteristics of the regulation of belief systems, the question arises of which of these two grounds the epistemic assessability of belief, that is, of whether belief is epistemically assessable (a) because truth is *what* belief is regulated for or (b) because evidentially is *how* belief is regulated. More specifically, there seem to be four possibilities here:

- (1) Truth Ground: Belief is epistemically assessable in virtue of being regulated for truth.
- (2) Evidence Ground: Belief is epistemically assessable in virtue of being evidentially regulated.
- (3) Conjunctive Ground: Belief is epistemically assessable in virtue of being both regulated for truth and regulated evidentially.
- (4) Disjunctive Ground: Belief is epistemically assessable in virtue of being either regulated for truth or regulated evidentially.

Why is belief epistemically assessable whereas mental states like supposition or imagination are not? According to (1), regulation for truth is what makes it the case that the former is epistemically assessable and the latter are not. According to (2), it is rather because imaginations and suppositions are not *evidentially* regulated that they are not epistemically assessable. One's daydream about vacationing on the beach is not sensitive to the evidence that one is actually quarantined in one's

tiny New York City apartment, because daydreams are not evidentially regulated. These are not the sorts of considerations that daydreams are sensitive to. And it is in virtue of this that they are not epistemically assessable. According to (3), it is precisely because belief is evidentially regulated for truth—regulated for just that and regulated just that way—that it is epistemically assessable. Imaginations and suppositions are doubly disqualified, on this view. Finally, according to (4), regulation for truth and evidential regulation are each *sufficient* for epistemic assessability. On this view, belief is doubly qualified as epistemically assessable (perhaps this is why belief is a *paradigm* of epistemic assessability).

We think of the choice among (1)-(4) as important future work in epistemology. We cannot settle such a foundational issue here, but would like to propose that either (2) or (4) must be right: epistemic assessability is grounded either in evidential regulation or in regulation-for-truth-or-evidential-regulation. What (2) and (4) have in common is that evidential regulation is *sufficient* for epistemic assessability. And the consideration we want to put forward in favor of this is simply that this is by far the most plausible way to diffuse the puzzle captured in our inconsistent triad.

For, as we have argued in sections I and V, emotions *do* seem to be evidentially regulated. Anger tends to be formed in response to evidence of the presence of offense, to be reinforced by additional evidence of the presence of offense, and to be extinguished by evidence against the presence of offense. What anger is regulated *for* is detecting offense, not truth. But the *way* in which it is regulated for detecting offense is evidential. And in virtue of this feature anger, like belief, is epistemically assessable. The same holds, *mutatis mutandis*, for fear, joy, and other emotion types. Similarly, fear is typically formed in response to evidence of danger and typically dissipates in response to the evidence being undermined or outweighed. What fear is regulated for is the presence of danger, because danger is the formal object of fear. Accordingly, what fear and anger are regulated for is tied up with what makes them the emotions they are—their formal object. Still, what makes them the emotions they are need not be what makes them epistemically assessable. What makes them epistemically assessable, we propose, is that they are *evidentially* regulated for whatever they are regulated for. The fact that this is (at least part of) the way they are regulated means that it makes sense to ask whether this token anger indeed fits the evidence, whether that token fear is well supported by the evidence, and so on. It means it makes sense to epistemically evaluate token emotional experiences.

So, this is our solution to the problem posed by the inconsistent triad from section III. We reject Epistemic Exclusivity, because we deny that it is (only) in virtue of having truth as its formal object that belief is subject to epistemic norms. Take the norm *believe that p only if your evidence supports p* (i.e., supports *p*'s being true). On our view, belief is subject to this norm not (only) because truth is what it is regulated for, but (also) because evidentially is how it is regulated for it. Therefore, it is not (just) in virtue of the fact that truth is the formal object of belief that beliefs are epistemically assessable; rather, it is (also) because beliefs are *evidentially* regulated. Emotions, too, are evidentially regulated. This is why the following norm applies to fear: *fear x only if your evidence supports x's being dangerous*. Indignation is subject to the following norm: *feel indignant about y only if your evidence supports y's being unjust*. And so emotions too are epistemically assessable, despite having different formal objects. Epistemic Exclusivity is false.

It *should* be pointed out that this approach to the inconsistent triad does leave some loose ends that would need to be addressed. There is in the position defended here something

unattractively opaque about the link between emotions' formal objects and their manner of regulation. When a mental state is aimed at truth, we can see why evidence would play a central role in how it is regulated. The link between the "what for" and the "how" of regulation is then *transparent*. But when a mental state is aimed at loss, or danger, or injustice, it is less transparent why it should be regulated evidentially. There is nothing about the nature of the "what for" that makes us see *why* evidence would play a role in the "how." A complete theory of epistemic assessability would have to address such issues.

What are the implications of this approach for the debate over the nature of emotion? We have argued against the cognitivist and perceptualist solutions to the problem posed in section III. Notice, however, that for the most part we did not argue against cognitivism and perceptualism *as such*, but rather against (a) the idea that cognitivism offers a satisfactory explanation of the epistemic assessability of emotion and (b) the idea that perceptualism offers a plausible way of denying the epistemic assessability of emotions. On our view, emotions *are* epistemically assessable, but not (only) in virtue of being, or involving as constituents, cognitive or doxastic states. This view tolerates certain versions of both cognitivism and perceptualism: versions of cognitivism which include, in addition to a doxastic constituent, an "affectively hot" constituent, and insist that this affective constituent is evidentially regulated, that is, that emotions are evidence-sensitive also in virtue of having this constituent; as well as versions of perceptualism that embrace a non-standard perceptual epistemology wherein perceptions are evidence-responsive.

There are, of course, many other theories of emotion, and for some of them our view may have more decisive implications. For instance, according to the so-called James-Lange theory of emotion, emotions are essentially bodily sensations (more precisely: proprioceptive perceptions of internal bodily events). Insofar as bodily sensations are paradigms of epistemic non-assessability, we find the James-Lange theory woefully inapt. Conversely, our view is quite congenial to primitivist or non-reductive approaches to emotion. Cognitivism, perceptualism, and the James-Lange theory are often put forward with a reductive ambition, namely, to reductively explain emotions in terms of mental states of other categories. But some philosophers have resisted any such reduction, claiming that emotions constitute a *sui generis* category of mental state, perhaps with an irreducibly affective phenomenology (Dewlaque 2017). The notion—defended here—that emotions are epistemically assessable but not in virtue of being (or involving) cognitive states, but rather in virtue of an evidence-responsive affectively hot component, is obviously congenial to such non-reductivism. To be sure, it does not quite *entail* emotion primitivism, nor does emotion primitivism entail *it*. Still, our account assigns to emotions a *sui generis* epistemic status, and this certainly *supports* the idea that emotions are a *sui generis* kind of mental state.

Conclusion

According to us, *all* mental states have a double intentionality: a particular object (often a proposition) they are directed at, and a formal object (always some value) they aim at. The formal object of a mental state type is what individuates it: what makes a mental state token belong to the type it does is that token's formal object. The formal object also imposes certain norms on the mental states, and dictates what the state is regulated for. In the case of belief, the formal object is truth, and this imposes truth-promoting norms on beliefs and dictates that what belief is regulated

for its truth. This may understandably suggest that belief is subject to epistemic evaluation *because* truth is its formal object, and as a corollary, that any mental state which does not have truth for a formal object is not subject to epistemic evaluation. However, in this paper we argued that examination of the epistemology and double intentionality of emotions shows this to be wrong. Emotions do not have truth for a formal object and yet they are subject to epistemic evaluation. We have proposed that what makes emotions epistemically evaluable is not *what* they are regulated for but *how* they are regulated for it – in particular, the fact that they are *evidentially* regulated is a sufficient ground for their epistemic evaluability. This proposal has far-reaching consequences for one's conception of *the sphere of the epistemic*. If we are right, mental phenomena concerned with truth form only a small subset of the epistemic domain. Any constitutively evidence-responsive mental phenomenon is subject to epistemic evaluation, regardless of its formal object.⁶

⁶ For comments on previous drafts, we are grateful to Fabrice Teroni and two anonymous referees for *Topoi*.

References

- Boghossian, P. A. 2003. 'The Normativity of Content.' *Philosophical Issues* 13: 31-45.
- Chudnoff, E. 2015. 'Phenomenal Contrast Arguments for Cognitive Phenomenology.' *Philosophy and Phenomenological Research* 91: 82-104.
- Clifford, W. K. 1877. 'The Ethics of Belief.' *Contemporary Review* 29: 289-309.
- Dewalque, A. 2017. 'Emotional Phenomenology: Toward a Nonreductive Analysis.' *Midwest Studies in Philosophy* 41: 27-40.
- Döring, S. A. 2015. 'What's Wrong with Recalcitrant Emotions? From Irrationality to Challenge of Agential Identity.' *Dialectica* 69: 381-402.
- Feldman, R. 2005. 'Epistemological Duties.' In P. K. Moser (ed.), *The Oxford Handbook of Epistemology*. Oxford and New York: OUP.
- Glüer, K. and Wikforss, Å. 2009. 'Against Content Normativity.' *Mind* 118: 31-70.
- Gordon, R. 1987. *The Structure of Emotion*. Cambridge: CUP.
- Helm, B. W. 2001. *Emotional Reason*. Cambridge: CUP.
- Kenny, A. 1963. *Action, Emotion, and Will*. London: Routledge & Kegan Paul.
- Kriegel, U. 2014. 'Towards a New Feeling Theory of Emotion.' *European Journal of Philosophy* 22: 420-442.
- Marks, J. 1982. 'A Theory of Emotion.' *Philosophical Studies* 42: 227-242.
- Montague, M. 2009. 'The Logic, Intentionality, and Phenomenology of Emotion.' *Philosophical Studies* 145: 171-192.
- Morreall, J. 1993. 'Fear Without Belief.' *Journal of Philosophy* 90: 359-366.
- Nussbaum, M. 2001. *Upheavals of Thought*. Cambridge: CUP.
- Nussbaum, M. 2004. 'Emotions as Judgments of Value and Importance.' In R. C. Solomon (ed.), *Thinking about Feeling*. Oxford: OUP.
- Salmela, M. 2006. 'True Emotions.' *Philosophical Quarterly* 56: 382-405.
- Schachter, S. and J. Singer 1962. 'Cognitive, Social, and Physiological Determinants of Emotional States.' *Psychological Review* 69: 379-399.
- Shah, N. and Velleman, D. J. 2005. 'Doxastic Deliberation.' *Philosophical Review* 114: 497-534.
- Siegel, S. 2017. *The Rationality of Perception*. Oxford and New York: OUP.
- Solomon, R. 1976. *The Passions*. New York: Doubleday.
- Velleman, D. 2000. *The Possibility of Practical Reason*. Oxford and New York: OUP.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford and New York: OUP.