

Deliberative Control and Eliminativism about Reasons for Emotions

Abstract

Are there are normative reasons to have—or refrain from having—certain emotions? The dominant view is that there are. I disagree. In this paper, I argue for Strong Eliminativism—the view that there are no reasons for emotions. My argument for this claim has two premises. The first premise is that there is a deliberative constraint on reasons: a reason for an agent to have an attitude must be able to feature in that agent’s deliberation to that attitude. My argument for this premise is that in order to have reasons for an attitude, we need to be able to exhibit some relevant form of control over this attitude, and this relevant form of control is deliberative control. The second premise is that no one can deliberate to any emotion. My argument for this premise turns on the claim that there is no deliberative question that is settled by forming (or giving up) an emotion. I contend that this is so due to the well-known phenomenon of recalcitrant emotions: for any deliberative question that can be settled, there is no guarantee that the relevant emotional state will follow (or be revised). Strong Eliminativism follows from these two premises.

Introduction

Are there are normative reasons to have – or refrain from having – certain emotions? The dominant view is that there are. Getting fired seems to be a reason to be sad. A friend wronging you seems to be a reason to be angry. Your trustworthiness seems to be a reason for your partner to not be jealous. Your partner’s infidelity seems to be a reason to be upset.

The dominant view is compelling; yet, I will argue, it is mistaken: there are no normative reasons for emotions. In my view, normative reasons don’t just favor certain actions and attitudes, but also play a crucial role in *deliberation*, enabling us to figure out what to do and what attitudes to form or abandon. However, I will argue that there’s an important respect in which emotions seem to be outside of our rational agency: emotions are not under our *deliberative control*. Given this and the tight

connection between reasons and deliberation, I maintain that there are no reasons for emotions. I will call this claim ‘Strong Eliminativism about Reasons for Emotion’ – or ‘Strong Eliminativism’.¹

My focus is on *normative* reasons, as opposed to explanatory, causal, or motivating reasons. Normative reasons (henceforth just ‘reasons’) for Φ -ing are facts or considerations that count in favor of Φ -ing. If there are sufficiently strong reasons for you to Φ , then you should Φ .

My argument for Strong Eliminativism has two premises:

1. A fact F is a reason for agent A to have (or refrain from having) an attitude X only if A can deliberate to X (or to refraining from having X) at least partly on the basis of F.
 2. No one can deliberate to any emotion.
- ∴ Therefore, there are no reasons for emotions.

I will argue for these premises in §1 and §2, respectively.

§1 Premise 1: The Deliberative Constraint

Deliberation, as I conceive of it, is the cluster of thought processes concerned with *settling a deliberative question*.² Settling a question consists in forming or refraining from having³ an attitude.

¹ Maguire (2018) also argues for this view, but on entirely different grounds.

² Hieronymi (2005).

³ As I use terms, refraining from having an attitude isn’t the mere lack or absence of an attitude; rather, refraining from having an attitude is a more committal kind of *stance*. Broome (2013) argues that we can only deliberate to beliefs (and other positive states) and not to their absences. Drucker (2022), however, argues that we can reason to any change in attitude that can be expressed as the conclusion of an argument, including *absences* of attitudes. My position is effectively a middle point between Broome and Drucker: I maintain only that we can deliberate to refraining from having an attitude *in the sense in which this involves a committal stance*.

Two commonly discussed types of deliberation are doxastic and practical deliberation. Doxastic deliberation is paradigmatically concerned with the deliberative question *whether P*. Settling the question *whether P* consists in forming a doxastic attitude toward P. If one settles the question by judging that P, one thereby comes to believe that P. If one settles the question by judging that not-P, one thereby comes to believe that not-P. Practical deliberation is concerned with a broad cluster of deliberative questions – e.g. *whether to Φ* , *when to Φ* , *how to Φ* , *what to do*, etc. – the settling of which consists in the formation of an *intention*. For simplicity, I'll focus on questions of the form *whether to Φ* . Settling the question *whether to Φ* consists in either deciding to Φ or deciding not to Φ . If one decides to Φ , one thereby comes to intend to Φ . If one decides not to Φ , one thereby comes to intend not to Φ .

At this point, it'll be helpful to introduce some technical terminology to mark some distinctions that may not be tracked in ordinary language. I'll say that one deliberates 'to' an attitude iff one settles a deliberative question *by* forming the attitude in question. One deliberates to a (dis)belief that P just in case one settles the question *whether P* by forming a (dis)belief that P. Conversely, I'll say one deliberates 'away from' an attitude iff one antecedently holds this attitude and then deliberates to giving up this attitude – as is the case, for instance, when one deliberates away from belief that P to either disbelief that P or suspension of judgment about *whether P*.

By contrast, there are some deliberative questions that seem to be *about* some attitude without being settled *by* forming or giving up this attitude. Consider, for instance, the deliberative question *whether believing that P is valuable*. Although this question is *about* the belief that P, this question is not settled *by* forming a belief that P, but instead by forming a (dis)belief that believing P is valuable. In this case, one would be deliberating to a belief *that believing that P is valuable*. I will use the term 'deliberating about an attitude' to refer to deliberation regarding these sorts of questions. In the example just given, one deliberates *about* – but not *to* – the belief that P.

Furthermore, I'll say that the attitudes that themselves settle a question – that is, attitudes that are such that settling a question itself consists in forming these attitudes – are the 'direct results' of the relevant deliberation. Beliefs and intentions are the direct results of doxastic and practical deliberation, respectively. By contrast, when an attitude or state merely *causally follows* from deliberation, without the settling of the relevant question itself consisting in forming the attitude/state in question, I'll call it an 'indirect result' of that deliberation. Indirect results of deliberation are not, strictly speaking, things that one deliberates *to*. The alleviation of a headache that results from taking aspirin is an indirect result of my *decision* to take aspirin, but it's inappropriate to say I deliberate *to* the absence of a headache.

Sometimes deliberating *about* an attitude can indirectly result in having this attitude. If I deliberate to the belief *that believing that P is rational*, then, if I go on to form the belief that P on the basis of this higher-order belief, the belief that P is the indirect result of my original deliberation. Similarly, if I deliberate to an *intention* to believe that P (or to get myself to do so), and if this results in actions that ultimately end in a belief that P, this belief is an indirect result of my original deliberation. I'll argue in §2 that, to whatever extent emotions ever result from deliberation, they only do so indirectly.

1.1 Clarifying the Deliberative Constraint

With that said, let's move on to the deliberative constraint:

DELIBERATIVE CONSTRAINT: A fact F is a reason for agent A to have (or refrain from having) an attitude X only if A can deliberate to X (or to refraining from having X) at least partly on the basis of F.

Let's clarify a few features of this constraint.

The deliberative constraint imposes a necessary condition on reasons; namely, reasons for having an attitude must be able to feature in an agent's deliberation to that attitude. This necessary condition is agent-specific; a non-agent-specific deliberative constraint, by contrast, would claim that reasons must be able to feature in *someone's* deliberation – say, that of an ideal agent.⁴ Furthermore, the deliberative constraint is a particular version of the more general view that one's normative reasons for Φ -ing must also be capable of being the reasons for which one Φ 's.⁵ The deliberative constraint is not a thesis about the constitution of reasons; as such, it should be distinguished from the reasoning-first view of reasons, the view that reasons-facts are explained by reasoning-facts.⁶

It should also be observed that the deliberative constraint seems to rule out the possibility of there being 'wrong-kind' reasons for attitudes. It's controversial exactly how to characterize wrong-kind reasons; to illustrate by example, they're the kind of reasons that are paradigmatically exemplified by pragmatic reasons for belief. Some philosophers hold that wrong-kind reasons for attitudes are genuine reasons for those attitudes, even though one apparently cannot deliberate to those attitudes on the basis of them.⁷ But the deliberative constraint entails that no such reasons exist: if one cannot deliberate to an attitude on the basis of some fact F, then F cannot be a reason for that attitude. It's worth noting a choice-point for theorizing here: some (e.g. Kolodny 2005; Shah 2006) argue from the truth of the deliberative constraint to the impossibility of wrong-kind reasons, whereas others (e.g. Howard 2019) argue from the possibility of wrong-kind reasons to the falsity of the deliberative

⁴ Even the non-agent-specific constraint would, together with premise two, support Strong Eliminativism. But I think that if one rejects the second premise, the agent-specific constraint would still have interesting normative upshots with respect to emotions.

⁵ The reasons for which one Φ s are sometimes called 'motivating' reasons.

⁶ See Setiya (2014) and Way (2017) for defenses of the reasoning-first view, and Schmidt (2020) for objections.

⁷ Leary (2017) argues that one can adopt certain attitudes on the basis of wrong-kind reasons, but even she doesn't say that we *deliberate to* attitudes on the basis of wrong-kind reasons.

constraint. I side with those who move in the former direction, but due to space constraints, I can't offer a defense of this choice here.

To say that one *can* deliberate to an attitude is to say that one has the *ability* to deliberate to an attitude. Though difficult to precisely characterize abilities, I'll offer a few remarks. Having an ability involves holding fixed an agent's psychological skills and mechanisms. However, certain obstacles, like opportunity costs, distractions, or a lack of motivation, do not constitute inabilities. Let's say an attitude is under one's *deliberative control* iff one can (in the just-defined sense of ability) deliberate to it. The deliberative constraint therefore says that an attitude needs to be under one's deliberative control in order for there to be reasons for having it.

1.2 *Why Accept the Deliberative Constraint?*

Many philosophers have defended some version of a deliberative constraint on reasons.⁸ In this section, I'll argue for my own preferred deliberative constraint, as stated above.⁹

There are reasons for some states but not for others. There are reasons for beliefs and intentions, but there aren't reasons for perceptions, headaches, and indigestion.¹⁰ We need some explanation for why there are reasons for some states but not for others. One plausible explanation of the difference is that we can exhibit a relevant form of control over our beliefs and intentions but not over perceptions, headaches, and indigestion. The principle undergirding this explanation is that in order for there to be reasons for an action, attitude, or state, we need to be able to exhibit some

⁸ For defenders, see Shah (2006), Kolodny (2005), Williams (1981), Searle (2001), Schroeder (2007: ch. 2), and Hieronymi (2005). For an objection, see Markovits (2014: ch. 2).

⁹ Similar arguments have been advanced by Kolodny (2005), Scanlon (1998), Gibbons (2010), and Way (2017).

¹⁰ Siegel (2017) argues that perceptions can be rational, but it's unclear whether she'd say that there are reasons for perceptions.

relevant form of control over that action, attitude, or state. Call this principle the ‘minimal control principle’.

To fill out the minimal control principle, we need to say what kind of control over some state is required to have reasons for it. After all, there is a *sense* in which we can control what we perceive – we can choose to open our eyes. But this form of control isn’t of the relevant form that would enable us to have reasons for perceptions. Similarly, there is some sense in which we can control our headaches or indigestion – we can take actions to prevent or soothe them. But this form of control isn’t the form relevant to normative reasons. Instead, we can have reasons to take actions to prevent or soothe headaches, because we have the relevant form of control over these actions.

In the case of reasons for action, the form of control philosophers paradigmatically identify is *voluntary* control. Simply put, voluntary control is the ability to do something at will. We have voluntary control over our actions, and – many have thought – that’s why we can have reasons for actions. But as has been argued by many others, it’s doubtful that we have voluntary control over our *attitudes*.¹¹ I can’t, for instance, will myself to believe that the number of stars is even or that government spending causes inflation. So, if we lack voluntary control over our attitudes, then we must either (a) identify some other relevant form of control we can exhibit over our attitudes, (b) give up the claim that we have reasons for any attitudes, or (c) give up the minimal control principle. Both (b) and (c) strike me as untenable, so I suggest we find another form of control.

More specifically, I suggest that the relevant form of control is deliberative control. Deliberative control is different from voluntary control; deliberative control is concerned not with an ability to do certain things at will, but rather with an ability to deliberate to particular states. We have deliberative control over beliefs and intentions: we can consciously consider reasons bearing on them in a way that counts as deliberating, and thereby come to believe and intend for those reasons. By

¹¹ See Williams (1973), Hieronymi (2006), McHugh (2012), and Helton (2018), among many others.

contrast, we can't deliberate to perceptions, headaches, and indigestion, and so we lack deliberative control over these states.

The deliberative constraint is effectively the product of the minimal control thesis and the idea that the relevant form of control, at least in the case of attitudes, is deliberative control. The deliberative constraint predicts precisely what we're presently trying to explain: we can have reasons for beliefs and intentions, but not for perceptions, headaches, and indigestion. And it explains this difference among what reasons there can be by appeal to differences in deliberative control. This explanatory power, I think, is a strong reason to accept the deliberative constraint.

One might object that deliberative control, and more generally any relevant form of control, is relevant only for the attribution of *responsibility*, but not for what reasons there are. As this objection goes, for there to be reasons for some action, attitude, or state, it doesn't matter whether we have control over this action, attitude, or state. Rather, some relevant form of control is needed only for the attribution of responsibility: it's appropriate to hold someone responsible only for what they have control over. But we can have reasons for Φ -ing even if we're not responsible for Φ -ing or failing to Φ .

I have two responses to this objection. First, if we deny that we need to have control over what we have reasons for – thus effectively abandoning the minimal control principle – then we lose the ability to explain why we can't have reasons for states like headaches, perceptions, and indigestion. Indeed, if we say that we can have reasons for states we have no control over – but not be responsible for them – then it's not clear why we shouldn't then give the same verdict for perceptions, headaches, and indigestion: there are reasons for these states, but we're just never responsible for having or failing to have them. But this seems like the wrong result: there don't seem to be reasons for perceptions, and there certainly aren't reasons for headaches and indigestion. Second, it's not clear what would be the point of positing 'reasons' that we can't even *in principle* be held responsible for failing to respond

to. After all, one of the driving thoughts behind why we don't have reasons for perceptions, headaches, and indigestion in the first place is that we're (virtually) never held responsible for our perceptions, headaches, or indigestion. If we force such a sharp wedge between what reasons there are and what we can in principle be held responsible for, it becomes hard to see why we're even thinking there are such reasons at all. Thus, I don't think this objection to the deliberative constraint is decisive.

§2 Premise 2: Emotions and Deliberation

The purpose of this section is to develop and defend premise 2, the claim that no one can deliberate to any emotions.

Let's begin by narrowing my focus. I will focus on emotions with propositional objects, as opposed to emotions with non-propositional objects (e.g. love for someone, anger with someone) and emotions without objects (e.g. general moods). Emotions with propositional objects are directed towards a proposition or state of affairs – e.g. anger *that Tom stole from me* – and are therefore propositional attitudes. I restrict my focus to emotions with propositional objects because they are, I think, the *best* candidate for emotions that we could have reasons for, since other propositional attitudes (like beliefs) are paradigm instances of things we have reasons for. Thus, I take emotions with propositional objects to be the *hard case* for me: if I can show that there are no reasons for emotions with propositional objects, then I'll effectively have shown that there are no reasons for any emotions at all.

I'll remain mostly neutral on what exactly emotions are.¹² One might be tempted to think that our theory of what emotions are affects whether we can deliberate to our emotions. For example, one might think that, if judgmentalism – the view that emotions are evaluative judgments – is true, then

¹² See Scarantino & de Sousa (2018) for an overview.

surely we can deliberate to them. However, I think this move gets things dialectically backwards: instead of our theory of emotions informing what features we think our emotions have, we should observe the various features of emotions and use these to guide our theorizing about what emotions are. Whether we can deliberate to emotions should inform our theory of what emotions are.

2.1 *Doxastic Questions Bearing On Emotions?*

Let's now turn our attention to the question of whether we can *in principle* deliberate to emotions. It's pretheoretically tempting to think that we can. After all, emotions sometimes seem to follow an episode of deliberation. I wonder whether I should feel mad, I judge that I should be, and anger follows.

However, this pretheoretically tempting line of thought doesn't show that we can deliberate to emotions. It only shows that forming and giving up emotions sometimes *follow* a deliberation, but this is consistent with the emotion merely following *indirectly*. Thus, even when an emotion follows in the wake of deliberation, that doesn't show that one deliberated *to* the emotion. Deliberating to (or away from) an attitude is a matter of settling a question, the settling of which consists in the formation (or abandonment) of that attitude. Thus, to show that we can deliberate to (or away from) emotions, we need to identify a deliberative question such that settling it consists in forming (or abandoning) an emotion.

But now observe a general difficulty: what deliberative question could there be, such that settling it would consist in an emotion?¹³ Suppose I'm afraid that I have to give a conference talk tomorrow. Being a pessimistic person, I believe, despite evidence to the contrary, that I'll do poorly

¹³ This difficulty is quickly raised by Sharadin (2016).

and, from this, that it's bad that I have to give a conference talk. But (let's suppose) I want to deliberate away from this fear. What deliberative questions could I ask myself? A few candidates jump to mind:

- (a) Whether I will do poorly at this conference talk
- (b) Whether it's bad that I have to give a conference talk
- (c) Whether I should be afraid that I have to give a conference talk
- (d) Whether giving a conference talk is something to be afraid of

These all seem like reasonable questions to ask oneself. However, the problem with (a)-(d) is that they are all *doxastic* questions, the settling of which consists in forming or abandoning certain *beliefs*, rather than certain emotions. Insofar as such questions bear on emotions, such questions seem to be cases of deliberation *about* emotions.

It might be suggested that it is enough for the view that we can deliberate to emotions if settling these doxastic questions is necessarily accompanied by the relevant emotional states. But it's evidently possible that all of these questions can be settled without the relevant emotional state following. That's because emotions sometimes don't align with our judgments.¹⁴ One way this happens is when an emotion is *recalcitrant* in the face of opposing judgments. An emotion is recalcitrant just in case the emotion *persists* despite the subject of the emotion making a judgment that is in tension with the emotion – for instance, remaining afraid that I have to give a conference talk even *after* I've judged, in light of my evidence, that the talk will go well, that fear is unfitting, and that I shouldn't be afraid.

¹⁴ One might object that, if, on my account, there are no reasons for emotions, then emotions can't *conflict* with a judgment at all. In response, I think we can make perfect sense of the phenomenon: an emotion and a judgment conflict just in case either (a) one has both E and a negative judgment regarding having E, or (b) one lacks E but has a positive judgment about having E. There may be a plurality of different contents involved in these positive and negative judgments. One might judge that E is unfitting, that it is bad, or that one shouldn't feel it. Indeed, one might even judge that one has decisive reasons against E: though this judgment, according to my view, must be false, this doesn't mean it can't be made.

Just as importantly, the converse phenomenon also seems to occur: sometimes we *don't* form an emotion despite making a judgment that would seem to call for the emotion – for instance, not feeling sad that my grandmother has passed away even after making a judgment that I should feel sad. For simplicity, I'll focus on the case of recalcitrant positive emotions, but everything I will say applies equally to the converse case as well.

It's a psychological fact that we experience recalcitrant emotions. The recalcitrance of emotions is such a widely experienced phenomenon that it's often treated as an essential data point for theorizing about emotions: if a theory can't explain how recalcitrant emotions are possible, then this is a decisive reason to reject the theory.¹⁵

Now consider again the doxastic questions (a)-(d). Cases of recalcitrant emotions show that our judgments bearing on emotions don't necessarily effectuate changes in our emotional states: we can and do make judgments that conflict with our emotional state. We should conclude from this the more general conclusion that whatever judgments we make bearing on our emotions, the appropriate emotion might not follow. That is, (a)-(d) could be settled, yet the emotion may not follow. I could judge that I won't do poorly at the conference talk, that the conference talk will be good for my career, that I shouldn't be afraid, and that a conference talk is nothing to be afraid of. Yet, my fear could persist even after making all these judgments – i.e. even after (a)-(d) have all been settled, and settled in ways that conflict with my fear. Therefore, settling (a)-(d) is not necessarily accompanied by the relevant emotional state.

¹⁵ D'Arms and Jacobson (2003), for instance, use the problem of recalcitrant emotions to argue that judgmentalism is false.

2.2 *A Sui Generis Emotional Deliberative Question?*

I argued in the last subsection that (a)-(d) are not settled by forming or abandoning an emotion. Is there any other deliberative question that is settled by forming or abandoning an emotion? Continuing again with the example of giving a conference talk, here is one potentially promising candidate deliberative question:

(e) Whether to be afraid that I have to give a conference talk

Let's say that (e) is an instance of the schematic question *whether to feel emotion E*.

Nevertheless, it's a little unclear how to understand this question. On one way of understanding this question, it seems to be a *practical* deliberative question – an instance of the schematic practical question *whether to Φ* where Φ is instantiated with *feel E*. But if that's how *whether to feel E* is understood, then this question is settled by forming an intention to (not) feel E, rather than by forming or abandoning E.¹⁶ So this interpretation of the question *whether to feel E* does not vindicate the proposal that it is a deliberative question that is settled by forming or abandoning an emotion.

Another way that the question *whether to feel E* can be understood is as reducing to one of the doxastic questions from the previous subsection. But again, this won't work for my opponent's purposes, because, as we just saw, these doxastic questions aren't settled by forming or abandoning emotions. Again, then, this interpretation of the question *whether to feel E* does not vindicate the proposal that it is a deliberative question that is settled by forming or abandoning an emotion.

¹⁶ Although we do sometimes form such intentions, there's something slightly unnatural about such intentions, because we don't ordinarily take our emotions to be under our voluntary control in the way that actions are. Instead, we more often ask ourselves practical questions about whether/how to *manage* our emotions.

What the proponent of this proposal needs, then, is the claim that *whether to feel E* is a sui generis “emotional” question that isn’t to be understood as either a practical or doxastic question. On this view, the question *whether to feel E* hasn’t been settled unless relevant the emotion has been formed.

We can try to flesh out this proposal by analogy to the question *whether to believe P*. *Whether to believe P* is not usually taken to be an instance of the practical question *whether to Φ*, where the latter is settled by an intention to Φ. Rather, *whether to believe P* is taken to be settled by forming a (dis)belief that P.¹⁷ Thus, it might be suggested that, analogously, the question of *whether to feel E* should be understood as a sui generis emotional question: just as settling the question *whether to believe P* issues directly in a belief, settling the question *whether to feel E* issues directly in an emotion.

I don’t deny that we can ask ourselves *whether to feel E* in thought. The issue, however, is whether this sui generis reading (according to which the question has only been settled if the relevant emotion has been formed or abandoned) is the correct interpretation of this question. I’ll argue that it is not.

The problem with the sui generis reading is that it’s prima facie possible that E (or its absence) can be recalcitrant even after the question *whether to feel E* has been settled. If there were a sui generis question, then settling it would necessarily issue in the relevant emotional state. But, I maintain, there is no deliberative question such that this is so. To see this, imagine that I’m asking myself whether to feel afraid that I’m giving a conference talk. I look at all the considerations that I take to bear on this question, and I’m *fully persuaded* that these considerations conclusively answer the question in the negative – i.e. support *not* being afraid. To all intents and purposes, it seems that I’ve settled this question. Yet, it seems plausible that, despite being so persuaded, my fear can still persist. Thus, it seems that the sui generis reading of *whether to feel E* is not the correct reading: *whether to feel E* can be

¹⁷ Proponents of transparency claim that *whether to believe P* simply reduces to the doxastic question *whether P* (Shah, 2003). But even those who reject transparency still accept that *whether to believe P* is settled by forming a (dis)belief that P; see [REDACTED].

settled without the relevant emotional state following, which means that *whether to feel E* cannot be a sui generis emotional deliberative question.

At this point, my opponent may respond that, appearances notwithstanding, you just don't count as settling the question *whether to feel E* until the relevant emotion has been formed or abandoned. In recalcitrant cases, the claim is, one *hasn't* yet settled the question *whether to feel E*. Rather, one might say, what it would be to settle this question *just would* be forming or giving up E. If one hasn't formed or given up E, then one hasn't really settled the question of whether to feel it.¹⁸

This suggestion can be bolstered by an analogy with weakness of will. If weakness of will is possible, then it may seem that there are cases where one can be persuaded by the relevant considerations one takes to bear on *whether to Φ*, yet the intention to (not) Φ does not follow.¹⁹ In such cases, it seems appropriate to say that one hasn't *settled whether to Φ* at all, precisely because one has not decided, and thereby come to intend, (not) to Φ. So, by analogy, perhaps recalcitrant emotions are much like cases of weakness of will: despite being so persuaded by the relevant considerations one takes to bear on the question, the question hasn't been settled at all. Surely, so says my opponent, we would not infer from cases of weakness of will that intentions don't settle *whether to Φ*; hence, nor should we infer from cases of recalcitrant emotions that emotions don't settle *whether to feel E*.

¹⁸ An alternative response to my argument on behalf of my opponent is that settling *whether to feel E* is a necessary but not sufficient condition for forming or abandoning E. However, this response seems to clearly deliver the wrong verdicts for more familiar cases of deliberation such as doxastic deliberation. If a deliberative question can be settled without thereby forming the corresponding attitude, this delivers the result that, in doxastic deliberation, one can settle the question *whether P* in the affirmative without thereby believing P. But "I've settled that P is true but don't believe P" sounds akin to a Moore-paradoxical statement. The lesson here is that, absent the relevant attitude, it's not clear how the question has been settled at all, nor what we'd even mean by saying the question has been settled.

¹⁹ It should be noted that it's not *obvious* that this is true in cases of weakness of will. Weakness of will is typically characterized – at least by those who do not distinguish weakness of will and akrasia – in terms of one's *believing that one ought to Φ* yet not intending to Φ. But it's not obvious that believing one ought to Φ constitutes being persuaded by the relevant considerations that one takes to bear on *whether to Φ*. I'll grant this for the sake of argument, however.

Holton (1999), by contrast, characterizes weakness of will in terms of intending to Φ without thereby bringing oneself to Φ. This *clearly* doesn't involve being persuaded by the considerations that one takes to bear on *whether to Φ* while not intending to Φ. On the contrary, it explicitly involves intending to Φ. So I assume this is not the notion of weakness of will my opponent has in mind.

My response to this suggestion, however, is that there are important and telling disanalogies between weakness of will and recalcitrant emotions.

It seems highly plausible that weakness of will can be explained by *uncooperative desires* – that is, uncooperative desires explain why one can be persuaded by the relevant considerations without the relevant intention following. Suppose I'm considering *whether to give to charity*, and I'm persuaded by the considerations I take to bear on this question in the affirmative. Indeed, I fully believe that I should give to charity. But now suppose further that I simply don't want to give to charity (or my desire to not give to charity is stronger than my desire to give to charity), even after considering the question of whether to do so. If I were to come to desire giving to charity (or if my desire to give to charity were to outweigh my desire not to do so), then I'd form the relevant intention. This is a paradigmatic instance of weakness of will: it is a case where I have judged that I ought to do some course of action, yet I don't thereby come to intend to do so. The foregoing considerations suggest that explaining weakness of will in terms of uncooperative desires is a powerful and plausible explanation. Though not all cases of weakness of will are exactly like the charity case – cases of depression, for example, seem importantly different – I suspect that even these other cases can ultimately be explained in terms of uncooperative desires.²⁰

By contrast, recalcitrant emotions can't be explained in terms of uncooperative desires. It's unnatural to say that someone experiencing a recalcitrant emotion *wants* to be experiencing that recalcitrant emotion. In fact, quite the contrary is normally the case: people typically want to *not* be experiencing a recalcitrant emotion.

²⁰ To elaborate a little: suppose I'm suffering from severe depression, and I'm considering the question *whether to get out of bed*. Suppose I'm fully persuaded by the considerations in the affirmative, and I fully believe that I should get out of bed. It seems right to say that depression messes with my motivational profile: because I'm depressed, either I don't want to get out of bed, or whatever desire I have to get out of bed is outweighed by an even stronger desire to stay in bed. Note that this is consistent with the possibility that there's a sense in which I'm *unable* to desire differently – and so, in this same sense, unable to intend to get out of bed – due to the severity of my depression. However, I am able to get out of bed in a different sense, namely: were my motivational profile to change, I would intend to get out of bed. As noted earlier, this is the sense of 'ability' at play in the deliberative constraint.

Why is this disanalogy relevant? It pumps an important pre-theoretical intuition about one's *abilities*: if the explanation for why I don't do something (or form some attitude) is that I don't want to do (form) it, then there's a good sense in which I'm *able* to do it. And correspondingly, if I *do* want to do something yet the action doesn't follow, then there's a good sense in which I'm *unable* to do it. Thus, if, in the case of recalcitrant emotions, I really want to stop feeling afraid but the fear won't go away, then it's natural to say that I'm unable – in the sense I've identified – to abandon my fear. Even though my desiderative profile is lined up in the right way, I still don't abandon the emotion, and hence am *unable* to do so in the sense I've identified. That is, it's because people usually *don't* want to experience recalcitrant emotions yet still do that they count as being unable to abandon their recalcitrant emotions.²¹

Why are abilities relevant? I propose that the notion of ability is crucial to understanding whether some attitudes settle a question. Call the attitudes that settle a question 'settling attitudes'. For any question Q, we can ask of any attitude X whether X is a settling attitude of Q. Something like the following, I suggest, is a general condition on whether X is a settling attitude of Q: if X is a settling attitude of Q, and if agent A is fully persuaded by the considerations A takes to bear on Q (in support of X), then A must be *able* to form X. Call this the 'settling condition'. Put another way: being persuaded by a set of considerations seems to be sufficient for settling a question, so if an attitude is a settling attitude of this question, then one must be in a position to form this attitude when one is persuaded by these considerations. If, despite being persuaded, one can't form this attitude, then this attitude can't be the settling attitude. For instance, take doxastic deliberation: again, in being persuaded

²¹ My opponent may concede that uncooperative desires aren't in one's way in cases of recalcitrant emotions but then argue that something *else* stands in one's way without constituting a genuine inability. For example, perhaps recalcitrant emotions can be explained by *intrusive thoughts*, and these don't constitute *inabilities* to feel (or not feel) some emotion. It's a little hard to assess whether this is what explains recalcitrant emotions without getting lost in the weeds of speculative psychology, but I'm skeptical that this explanation can really generalize to *all* recalcitrant emotions.

by the considerations I positively take to bear on the question *whether P*, I simply come to believe P. Thus, believing *P* meets the settling condition for the question *whether P*.

Emotions, however, fail the settling condition for the question *whether to feel E*. I submit that, when one finds the considerations bearing on *whether to feel E* fully persuasive, yet the relevant emotional state doesn't follow, one is genuinely *unable* to form the relevant emotion. Therefore, given the settling condition, some other attitude must be the settling attitude for this question.

By contrast, intentions – even considering cases of weakness of will – do not seem to fail the settling condition. Recall the sense of 'ability' that is relevant for whether one *can* deliberate to an attitude in the sense at play in the deliberative constraint. Uncooperative desires don't constitute inabilities in this sense of the term. So, if weakness of will is explained in terms of uncooperative desires, then weakness of will doesn't involve a genuine *inability* to form the relevant intention: if one were to desire differently, then one would form the relevant intention. Although it might be *difficult* to so intend, uncooperative desires don't constitute inabilities in my sense of the term.

To sum up this complex and subtle dialectic, the proposal on the table was that the question *whether to feel E* is a sui generis emotional deliberative question, which is settled by forming or abandoning E. In response, I suggested that the question *whether to feel E* can be settled without the relevant emotional state following and that, therefore, *whether to feel E* should not be given this sui generis interpretation. In response to this, my opponent suggested that in such cases one hasn't *really* settled the question at all. I responded to *this* by arguing that there is a plausible condition – the settling condition – according to which if attitude X is a settling attitude of Q, and if agent A is fully persuaded by the considerations A takes to bear on Q (in support of X), then A must be able to form X. In cases of recalcitrant emotions, I argued, A is unable to form (or give up) E. Thus, emotions fail the settling condition for the question *whether to feel E*. Therefore, settling this question consists in some other attitude than an emotion. Therefore, this question is not a sui generis emotional deliberative question;

instead, I think, it's (most commonly) shorthand for one of the doxastic questions considered in the last subsection – such as the question *whether I should feel E* – which are settled by forming beliefs.

2.3 *Completing the Argument: All Emotions Are Not Under Our Deliberative Control*

I have argued in the previous subsections that from the phenomenon of recalcitrant emotions, we can see that there is no deliberative question such that settling it consists in forming (or giving up) an emotion. But it bears stressing that this does not just show that we can't deliberate to (or away from) emotions *in recalcitrant cases*. Rather, it shows that we can't deliberate to (or away from) emotions more generally. Even if emotions do sometimes follow the settling of the question (in non-recalcitrant cases), the settling of the question can't *constitute* the forming or abandoning of the emotion (due to recalcitrant cases). Therefore, in settling the question, we're not deliberating *to* (or away from) the relevant emotion, even in non-recalcitrant cases. Rather, we're deliberating *about* the emotion, and in non-recalcitrant cases forming (or refraining from having) the emotion follows indirectly. If my opponent were to grant that we can't deliberate away from non-recalcitrant emotions, yet insist that we can deliberate to (or away from) non-recalcitrant emotions, they'd need to argue that there's a relevant deliberative question *only* in the non-recalcitrant cases. But if there were such a question, then settling it would always consist in forming an emotion. But this applies equally to recalcitrant emotions: settling this question would be able to unseat recalcitrant emotions. So, there couldn't be such a question *only* in non-recalcitrant cases.

I therefore conclude that there is no deliberative question such that settling it consists in an emotion. Rather, settling all the relevant deliberative questions which bear on emotions consists in forming either beliefs or intentions, and the problem of recalcitrant emotions show that even after settling these questions, the corresponding emotional state may not follow. I therefore conclude that

emotions are not under our deliberative control – i.e. it’s not possible to deliberate to (or away from) emotions.²²

Thus, I maintain that we can only deliberate about, rather than to, emotions. Often, our hope is that these deliberations will change our emotions. Sometimes they do, and other times they don’t. But when they do, they only do so *indirectly*. We deliberate to a judgment that bears on an emotion, and this judgment is the direct result of the deliberation. This, then, leads to a change in one’s emotion. So, the deliberation at best indirectly led to a change in one’s emotion. The emotion changed *via* a change in belief. Hence, settling a question is *causally* but not *constitutively* tied to forming or abandoning emotions.

Our emotions are often sensitive to our judgments. However, this sensitivity is not under our rational agency, construed as a relation of deliberative control. It’s often very easy to feel the appropriate responses to these judgments; in a huge swath of ordinary cases, nothing stands in the way between our judgments and fittingly feeling emotions in response to those judgments. When that happens, our emotional faculties are working ‘properly’, so to speak.

But many other times, it’s quite hard for people to feel the appropriate responses to the various judgments and/or situation. Some emotions are harder to *rid* ourselves of than others (think of recalcitrant emotions). Some emotions are harder to *feel* than others (think of a person who *wants* to feel sad at a funeral, but just can’t). Some emotions are harder for *me* to feel (or rid myself of) than they are for *you*. The ease and difficulty of having or ridding ourselves of certain emotions comes in

²² It’s important that this argument doesn’t overgeneralize to rule out deliberating to beliefs. One might worry that if the recalcitrance of emotions is supposed to show that we can’t deliberate to emotions, then the recalcitrance of *belief* shows that we can’t deliberate to beliefs; this, so says my opponent, would be an absurd result. Drucker (2022) raises this sort of worry. My response is that the recalcitrance of belief and emotion are importantly different. Even if there is some sense in which belief can be recalcitrant - for example, when one continues to believe *P* in the face of a higher-order judgment that one has insufficient evidence for *P* - what is *not* possible is to have settled the question *whether P* in the affirmative and yet to not believe *P*. If one has settled the question *whether P* in the affirmative, one *thereby* believes *P*. Thus, even if belief can be recalcitrant in the face of a higher-order judgment, there is *a* deliberative question (namely, *whether P*) that just settles belief, such that belief *cannot* be recalcitrant in the face of having settled it. By contrast, I am arguing, there is no such deliberative question in the case of emotion.

degrees and varies across emotions, situations, and people. For many people who experience persistent negative emotions like shame or self-loathing, it can be so tough to get rid of these emotions that they determine (or are encouraged by others) that they should undergo psychotherapy to work through the difficulty and effectively retrain their emotional responses.

In sum, sometimes deliberation leads to a change in emotion, but this relation is indirect rather than direct. Our emotional faculties are working properly just in case our emotions are fitted to the relevant background judgments that are associated with our emotions, yet even when our emotional faculties are working properly, our emotions are still not under our deliberative control. It's not possible to deliberate to, or away from, emotions.

Final Remarks

In this paper, I've argued for Strong Eliminativism – the view that there are no reasons for emotions. One might worry, however, that this conclusion is too revisionary to be plausible, and so my argument must have gone wrong somewhere. After all, intuitively there are reasons for emotions, and this is reflected in how we talk about emotions. Relatedly, if there are no reasons for emotion, then we might worry that we lose the ability to morally criticize people for their intuitively unwarranted and even pernicious emotions. This is, in effect, a kind of 'Too Few Reasons' objection: if a theory of reasons rules out the existence of reasons in cases in which we ordinarily take there to be such reasons, then, *ceteris paribus*, we should reject the theory.²³

Think of the jealous boyfriend. The jealous boyfriend lacks a reason to believe his partner is untrustworthy, and his partner continuously proves their honesty through their actions, yet the

²³ The Too Few Reasons objection has been discussed in response to a variety of theories about reasons, including moral subjectivism, moral error theory, and epistemic instrumentalism. See, e.g., Schroeder (2007), Côté-Bouchard (2015).

boyfriend gets jealous whenever his partner hangs out with one of their attractive friends. Even worse, the boyfriend becomes possessive and tries to undermine his partner's agency. Intuitively, the jealous boyfriend has reasons to not be jealous, and this is reflected in what we might say, e.g. 'his partner has given him every reason to not be jealous'. Moreover, denying that he lacks reasons to not be jealous seems to prevent an avenue for criticizing his immoral behavior.

I accept that we have intuitions about the appropriateness of emotions, and that we sometimes express these intuitions using the word 'reason'. The question, however, is whether this usage really maps onto the specific notion of a reason that is at work in philosophical theorizing, and whether we really need to invoke this notion to capture these intuitions. I think that we can accommodate our intuitions about the appropriateness of emotions without positing reasons for emotion. First, we can capture some such intuitions by appeal to reasons for other states, such as beliefs: the jealous boyfriend has sufficiently good reasons to *believe* that his partner is trustworthy. Second, we can capture others by saying that emotions are *fitting* or *unfitting*. For example, because the jealous boyfriend has sufficiently good reasons to believe that his partner is trustworthy, jealousy is therefore unfitting. My account does not preclude this possibility; rather, it merely precludes the view that such fittingness-facts (always) constitute normative reasons.²⁴ Perhaps these ways of capturing our intuitions would make utterances like 'his partner has given him every reason to not be jealous' false at face value. But we can accept a plausible linguistic error theory whereby there are true utterances (say, about fittingness) *in the neighborhood* of reasons-talk. Therefore, I think we can adequately account for our intuitions and utterances without positing reasons for emotions.

²⁴ One might be puzzled that I deny that there are normative reasons for emotions yet concede that emotions can be fitting. On a simple fittingness-first view of normative reasons, a fact F is a reason for an attitude or action iff, and because, F makes the attitude or action fitting. On this view, if emotions can be fitting, then there are reasons for them. However, if we deny the simple version of the fittingness-first view, then we can allow that emotions can be fitting without conceding that there are reasons for emotions. My view is that we should deny at least the *simple* version of the fittingness-first view, precisely because it commits us to the existence of reasons that (I have argued) fail the deliberative constraint. My view, however, is compatible with other versions of fittingness-first views. Strong Eliminativism merely constrains the acceptable fittingness-first views we can accept. See Howard (2019) for a defense of fittingness-first views.

Moreover, in any case, it strikes me as obvious that we have more than enough resources to condemn the jealous boyfriend without appealing to reasons. Actions can be moral or immoral, possessive behavior is immoral, and this doesn't depend on there being reasons for emotions. So, if what we're concerned about are the actions stemming from an emotion, we can condemn the boyfriend for *acting on* his jealousy in immoral ways. It's not necessary to add a further criticism on top. If, on the other hand, what we're concerned with is condemning the emotion *itself*, then I think we can do so simply on the basis of the unfittingness of the emotion. As above, this doesn't require saying that there are reasons for or against the emotion. Therefore, I think we have all the resources we need to get the right moral conclusions.²⁵

At bottom, I'm inclined to think that many ways we philosophers theorize about normative reasons tend to obscure why reasons matter to us and what roles they play in our lives. By centering the way that reasons figure in deliberation, we regain a sense of why reasons matter to us and what roles they play: reasons help us figure out what to do and, accordingly, help us assign praise and blame. While my view may initially seem to underpredict the number and kinds of reasons we have, relative to what philosophers have traditionally thought, I think that because so many discussions of reasons displaced reasons from these central roles, philosophers hypothesized that we have reasons for a greater range of states than those which we really do. Indeed, it remains to be seen how my argument extends to other attitudes, say, preferences, desires, or even credences. It may turn out that beliefs and intentions are the only attitudes under our deliberative control. For some, this might just be more reason to give up the deliberative constraint. But I think that the central roles that reasons play in our normative lives give us strong reasons to focus our theorizing about reasons on states that we have the requisite control over. We seem to lack the requisite control over emotions, so I conclude that we lack normative reasons for them.

²⁵ Williams (1981) deploys a similar strategy in defense of his reasons internalism.

References

- Broome, J. (2013). *Rationality Through Reasoning*. Malden, MA: Wiley-Blackwell.
- Côté-Bouchard, C. (2015). Epistemic Instrumentalism and the Too Few Reasons Objection. *International Journal of Philosophical Studies*, 23(3), 337-355.
- D'Arms, J., & Jacobson, D. (2003). The significance of recalcitrant emotion (or, anti-quasijudgmentalism). *Royal Institute of Philosophy Supplement*, 52, 127-145.
- Drucker, D. (2022). Reasoning beyond belief acquisition. *Noûs*, 56(2), 416-442.
- Gibbons, J. (2010). Things That Make Things Reasonable. *Philosophy and Phenomenological Research*, 81(2), 335-361.
- Helton, G. (2018). If You Can't Change What You Believe, You Don't Believe It. *Noûs*, 54, 501-526.
- Hieronymi, P. (2005). The Wrong Kind of Reason: *Journal of Philosophy*, 102(9), 437-457.
- . (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87(1), 45-74.
- Holton, R. (1999). Intention and Weakness of Will. *Journal of Philosophy*, 96(5), 241-262.
- Howard, C. (2019). The Fundamentality of Fit. *Oxford Studies in Metaethics*, 14.
- Kolodny, N. (2005). Why Be Rational? *Mind*, 114(455), 509-563.
- Leary, S. (2017). In Defense of Practical Reasons for Belief. *Australasian Journal of Philosophy*, 95(3), 529-542.
- Maguire, B. (2018). There Are No Reasons for Affective Attitudes. *Mind*, 127(507), 779-805.
- Markovits, J. (2014). *Moral Reason*. Oxford: Oxford University Press.
- McHugh, C. (2012). Epistemic Deontology and Voluntariness. *Erkenntnis*, 77, 65-94.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scarantino, A. & de Sousa, R. (2018). Emotions. *Stanford Encyclopedia of Philosophy*.
- Schmidt, E. (2020). Where Reasons and Reasoning Come Apart. *Noûs*, 1-20.

- Schroeder, M. (2007). *Slaves of the Passions*. New York: Oxford University Press.
- Searle, J. (2001). *Rationality in Action*. Cambridge, MA: MIT Press.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies*, 167(2), 221-235.
- Shah, N. (2003). How Truth Governs Belief. *The Philosophical Review*, 112(4), 447-482.
- . (2006). A New Argument for Evidentialism. *The Philosophical Quarterly*, 56(225), 481-498.
- Sharadin, N. (2016). Reasons Wrong and Right. *Pacific Philosophical Quarterly*, 97(3), 371-399.
- Siegel, S. (2017). *The Rationality of Perception*. Oxford: Oxford University Press.
- Way, J. (2017) Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly*, 98, 251– 270.
- Williams, B. (1973). Deciding to believe. In *Problems of the Self: Philosophical Papers 1956–1972* (pp. 136-151). Cambridge: Cambridge University Press.
- . (1981). Internal and External Reasons. In *Moral Luck*, 101-113. Cambridge: Cambridge University Press.