

Why Fittingness Is Only Sometimes Demand-Like

James Fritz

This preprint has not undergone peer review or any post-submission improvements or corrections. The Version of Record of this article is published in Philosophical Studies, and is available online at <https://doi.org/10.1007/s11098-023-01971-1>.

Abstract Sometimes, the fact that an attitude is fitting seems like a demand to have that attitude. But in other cases, the fact that an attitude is fitting seems more like a permission to have the attitude. I defend a proposal that can accommodate both of these appearances. I argue that there is a kind of emotionlessness, which I call *apathy*, that can be fitting or unfitting in just the same way that emotion can. I further argue that, in some cases, it can be fitting to respond a single object either with emotion or with apathy. When both apathy and emotion are fitting options, the fittingness of the emotion is a permission-like status; failures to have the fitting emotion are not failures of fit. But when an emotion is fitting and apathy is unfitting, the fittingness of the emotion is a demand-like status; failures to have the emotion are failures of fit.

We can evaluate attitudes against many different standards. Take, for example, the enjoyment I feel while watching the movie *Point Break*. There are a variety of questions that I could ask about that enjoyment. Is it *morally vicious* for me to enjoy *Point Break*? Is it *healthy* for me to do so? Is it the *best response, all things considered*, that I could have to the movie? Each of these questions measures my enjoyment against a different standard for success.

Many hold that we can also measure my enjoyment of *Point Break* against a further standard: we can ask whether that enjoyment is *fitting*. Whether we can give an informative, non-circular definition of fittingness is a matter of controversy. But fittingness is a familiar property: it's fitting to feel shame when you've done something shameful, to admire people who are admirable, to be disgusted by meals that are disgusting, and to enjoy action movies that are enjoyable. To say that an attitude is fitting is to say that the attitude *matches its object* in a certain way.¹

Sometimes, the fact that an attitude is fitting seems to amount to a kind of *demand* to have that attitude. Consider, for instance, the notorious "Pharma Bro" Martin Shkreli, who engaged in egregious price-gouging while selling life-saving medicine. Shkreli's behavior was shameful; it was fitting (and it remains fitting) for him to feel shame. This fittingness-fact seems to be a fact that positively calls for Shkreli's response.

¹ For introductions to fittingness, see D'Arms and Jacobson (2000) and Howard (2018).

When we regard Shkreli's behavior as shameful, we are not simply noting that shame is an option for him. We are considering shame to be, in some sense, the *only* eligible option for him. If Shkreli fails to feel any shame, his response falls short. Some philosophers, impressed by cases of this sort, claim that fittingness is a demand-like status.

Sometimes, however, fittingness does not seem demand-like. Suppose that, at a yard sale, I notice that a cupboard's sleek door is beautiful—it makes aesthetic pleasure fitting.² Does this mean that, if I took no aesthetic pleasure in the cupboard door, there would be some problem with my response? Many will suspect not. Some philosophers, impressed by cases of this sort, take the fittingness of an attitude to be a permission-like, not a demand-like, status.

In this paper, I defend a proposal that can accommodate both of these inclinations; on my proposal, fittingness is sometimes demand-like and sometimes permission-like. The correct explanation for this diversity, I hold, emerges naturally when we shift our focus from questions about the fittingness of coarse-grained attitudes (like shame) to questions about the fittingness of fine-grained attitudes (like degrees of shame).

Section 1: A Riddle about Fittingness

1.1 Fittingness, Requirement, and Permission

The literature on fitting attitudes has, historically, treated fittingness as a kind of requirement.³ Discussions of the “buck-passing” account of value, for instance, frequently treat the following two proposals as interchangeable:

F is valuable if and only if it is *fitting* to value F.

F is valuable if and only if one *ought* to value F.⁴

I doubt that this tendency to conflate fittingness with requirement is merely a historical accident; the proposal that fittingness is a kind of demand is a natural first assumption. For one, there often seems to be a kind of problem, or a falling-short, when a person fails to have a fitting emotion. What's more, it's natural to use claims about the fittingness of emotion to call attention to the relevant problem. A person who learns that Martin Shkreli feels no shame about his misdeeds could criticize that lack of shame by saying, “what you did was shameful!” And a person whose joke fails to amuse a listener could criticize the listener by saying, “oh, come on; that joke's *funny*.” These claims sound like criticisms, and the view that fittingness is a kind of

² The cupboard example appears in Whiting (2021: 414).

³ [Redacted for blind review]

⁴ Ewing, for example, writes “If we mean by ‘good’ what ought to be desired, approved, or admired, it seems still more obvious to me that we are thinking of ‘ought’ in the sense in which it signifies fittingness” (2012, 151). See also Rabinowicz and Rønnow-Rasmussen (2004) and Gertken and Kiesewetter (2017).

requirement explains why: to call something shameful is to say that it's fitting to feel shame about it, and the fittingness of shame is a kind of requirement to feel it.

But the view that fittingness is a kind of requirement faces *prima facie* difficulties. In many cases, failures to have fitting emotions do not seem problematic in any sense. When I see a somewhat-beautiful cupboard door, but do not react with aesthetic pleasure, it's far from clear that anything has gone wrong. Nor is there any obvious problem when a person fails to be annoyed by a slightly-annoying car alarm or when a person does not worry after hearing somewhat-worrisome news about the stock market. When we as theorists consider cases like this, it's natural for us to have a *lenient* reaction—that is, to think that the people we're imagining aren't going wrong or falling short in any sense. This lenient reaction is difficult to square with the view that fittingness is a kind of requirement.

What else could fittingness be? One proposal looks to a different deontic category; it says that the fittingness of an attitude is always a kind of *permission* to have the attitude.⁵ But, whereas the view that fittingness is a kind of requirement seemed inadequately lenient, the view that fittingness is a kind of permission seems *too* lenient. On this view, the fact that Martin Shkreli's behavior is shameful would simply be the fact that he is *permitted* to feel shame. Likewise, the fact that a raging wildfire near one's home is fearsome would simply be the fact that one is *permitted* to fear the wildfire. This seems far too weak. Fear is not simply an *unproblematic* reaction to the wildfire. It is, in at least some sense, a *correct* response; there is something that fear of the wildfire *gets right*. The proposal that fittingness is nothing more than a kind of permission seems like a bad match for cases like this.

We have the makings of a riddle here. In some cases, the fittingness of an attitude seems more like a requirement than a permission. But in other cases, the fittingness of an attitude seems more like a permission than a requirement. What explains this variation?

One natural response to this riddle is to give up on the project of attempting to reduce fittingness to, or analyze it in terms of, deontic categories. This is Selim Berker's response; he marshals the data I've considered (and a great deal of other data as well) in order to argue that fittingness is a distinctive normative category—one that cannot be reduced to, or analyzed in terms of, either requirement or permission.⁶ On Berker's proposal, to say that an attitude is fitting is always to say something more than that the attitude is merely permitted. But it is also never to commit oneself to the attitude's being required.

Berker's proposal may well be correct. But, even if it is, it leaves our riddle about fittingness intact. The riddle, as I see it, is one about failures to have emotions that would be fitting. Put bluntly, there are some cases in which those failures seem like

⁵ Daniel Whiting defends a limited version of this conclusion, writing that, at least “in the aesthetic domain, an affect's being fitting is a permissive matter, not an obligatory one” (2021, 413-4).

⁶ Berker also argues against the proposal that the fitting can be understood in terms of evaluative categories, like *good* or *better than*.

an issue, but there are other cases in which those failures seem entirely unproblematic. Berker leans on precisely these intuitions at different stages of his argument against analyzing fittingness in deontic terms; he is, then, committed to taking these intuitions to be intuitions about fittingness, and to taking them seriously. But, even if Berker's conclusion is right—even if fittingness is a distinctive, non-deontic normative category—that point does not by itself illuminate the diversity in the intuitions to which Berker appeals. Berker's core idea, after all, is that we should lump all instances of fittingness together into a single (distinctive, primitive) class. And that move is manifestly not the sort of move that could explain why fittingness seems to behave differently in some cases than in others.

In other words: even if we grant that fittingness is neither a requirement nor a permission, we should acknowledge that there are some cases where it seems demand-like (in the sense that the absence of emotion seems problematic) and others where it seems more permission-like (in the sense that the absence of emotion is unproblematic). This variation cries out for explanation, and the proposal that fittingness is neither a requirement nor a demand does not, by itself, suffice to provide the needed explanation.

In what follows, I defend a model that does successfully explain how fittingness could be permission-like in some cases and demand-like in others. The model I defend, importantly, is neutral as to whether fittingness can be reduced to or analyzed in terms of other normative categories, like *demand*, *permission*, *goodness*, *badness*, *rightness*, *wrongness*, and *most reason*. I aim to show that we don't need to settle these questions about the relationships between normative categories to explain how fittingness could sometimes be demand-like and permission-like. We can provide that explanation with only a minimal set of tools: all we need is the property of fittingness and the property of unfittingness. Everyone, no matter their background commitments about the nature of fittingness or its relationship to other normative categories, has these tools in their toolkit. So everyone can comfortably embrace the explanation I go on to offer.

Here's a brief preview of that explanation. In the next section, I argue that, just as it can be either fitting or unfitting to *have* an attitude, it can also be fitting or unfitting to *lack* an attitude. This, I argue, gives us all the machinery that we need to explain why fittingness is sometimes demand-like and sometimes permission-like. When it's fitting to have an attitude and unfitting to lack the attitude, fittingness behaves like a demand, in the sense that failures to have the attitude amount to failures of fit. But when there are both fitting ways of having the attitude and fitting ways of lacking the attitude, fittingness behaves like a permission; failures to have the attitude in question needn't involve failures of fit.

1.2 The Status of the Riddle

Before I move on to defend this positive proposal, however, I'll consider some concerns about the riddle I've articulated. I've motivated that riddle by appealing to intuitions about the status of absent emotion in particular cases. But some might

suspect that those intuitions are not (or, at least, are not clearly) intuitions about fittingness. Perhaps, for instance, we have the sense that Martin Shkreli's lack of shame is problematic not because it is *unfitting* but because it manifests a *moral* failing—a moral failing that, notably, does not arise when one feels no aesthetic pleasure about a beautiful cupboard door.

There are two importantly different ways to press this worry. The bolder strategy involves identifying some *particular* non-fittingness standard—for instance, the standard of morally virtuous emotion—and claiming that our intuitions about whether fittingness is demand-like are systematically sensitive to that standard. This strategy takes our intuitive reactions to cases of absent emotion as important data points, and it attempts to explain those data points without appealing to fittingness.

This first strategy, however, relies on an implausibly simple story about our intuitions. It's not true, for instance, that fittingness only seems demand-like in cases of *moral* emotion; as I've already mentioned, we can call on a person to be amused by saying "oh, come on, that joke's *funny*." Other proposals in this vein seem equally unpromising; we shouldn't assume, for instance, that intuitions about absent emotion are all covertly intuitions about whether one has *most all-things-considered reason* to have the emotion, or whether doing so would have significant *non-instrumental value*.⁷ It's entirely possible to consider Martin Shkreli's lack of shame an *unfitting* reaction while also being quite convinced that he does not have *most all-things-considered reason* to feel shame, or while entirely convinced that shame never has non-instrumental value. The problem, in a nutshell, is that this too-bold debunking strategy sells us short as evaluators of absent emotion. It assumes that our intuitive reactions to absent emotion are always straightforwardly dominated by a particular non-fittingness standard. But our intuitive reactions to cases of absent emotion are more nuanced than that.

A more cautious debunking strategy is more convincing. This strategy, rather than taking our intuitive reactions to particular cases seriously, claims that those intuitions are too murky and unreliable to be trusted. When we have the intuition that an emotion goes wrong, after all, it can be difficult to tell whether we're reacting to a problem with the *fittingness* of emotion, or a problem with that emotion meeting some other standard. This is a perennial problem for theorizing about the fittingness-conditions for emotion, one that Justin D'Arms and Daniel Jacobson memorably dub "the opacity of normative force" (2014). Given this opacity regarding the norms picked out by our intuitions, the argument goes, we should downplay the importance of finding a theory that respects intuitions about cases. Perhaps we should prioritize finding a theory of fitting emotion that is elegant, simple, and explanatorily powerful—even if that theory does not vindicate all of our intuitions about cases.

I think that this second strategy is onto something important. Our intuitions about particular cases of absent emotion are indeed far from infallible guides to fittingness. And appreciating this fact should indeed drive us to downplay the importance of

⁷ See Rowland (2022) for a view on which questions about the authoritative normativity of fittingness are closely connected to questions about non-instrumental value.

finding a theory that respects all of our intuitions about cases. So, if the only way to vindicate our intuitions about the particular cases I've mentioned were to embrace a gerrymandered, unprincipled, or otherwise unattractive theory, we should be willing to give up on the attempt to vindicate those intuitions.⁸

Happily, however, we don't have to make that sacrifice. There is an entirely principled, independently attractive theory that nicely resolves our core riddle. It explains why fittingness is only sometimes demand-like, and it does so by appealing only to some weak and plausible assumptions. Even those who embrace a healthy measure of skepticism about our intuitions about fitting emotion, then, have excellent reason to take this theory seriously. In the next section, I introduce and defend that theory. And in section 3, I show how the theory can be used to resolve our core riddle.

Section 2: Fitting Emotion and Fitting Apathy

Many attitudes can be held with varying degrees of intensity. I felt happy the last time I had a good beer, and I also felt happy the first time that I saw my child smile—but the latter episode involved more happiness than the former. What's more, an attitude's degree of intensity can make a difference to whether it is fitting. It's fitting for me to be glad that I'll be having a crisp apple later today, but if I were absolutely ecstatic about it, my reaction would be unfitting. The fact that gladness is a fitting response to some object, then, does not mean that any degree of gladness will do.

In this section, I defend three theses about the relationship between fitting coarse-grained attitudes (like gladness) and fitting fine-grained attitudes (like degrees of gladness). These three theses, when jointly accepted, will put us in a position to explain why fittingness is only sometimes a demanding property.

Throughout the discussion that follows, I limit my focus to emotion. My goal is to show how it could be that fittingness is a demand-like property in some cases, and a permission-like property in others. If I prove that there are cases of both types when it comes to fitting *emotion*, that will suffice to prove the more general point about fitting attitudes. I suspect that the model I offer below will generalize nicely to show that the fittingness of some other propositional attitudes (including belief and hope) can also be either demand-like or permission-like. But I'll leave discussion of those other attitudes for future work.

2.1 Against Uniqueness about Fitting Degrees of Emotion

My first thesis is the denial of the following claim:

⁸ [Redacted for blind review.] Note that one can safely set aside some lenient intuitions while also taking seriously the lenient intuitions that I've marshalled in this section. This is precisely what Berker does; he notes that, although we should not be moved to leniency by concerns about our cognitive limitations, there are good grounds for favoring a lenient approach that have nothing to do with cognitive limitations. "When someone tells a cringeworthy joke," he writes, "it is not compulsory for me to cringe, but that is not because cringing would expend crucial emotional resources that could be devoted elsewhere" (forthcoming, sec. 4).

Uniqueness about Fitting Degrees of Emotion Any time an object merits an emotional reaction from a person, there is at most one degree of intensity to which it's fitting for the person to have the relevant emotional reaction.

I think it's very plausible that this claim is false—that, in other words, there are at least some cases in which there is no unique degree of fitting intensity for a (coarse-grained) emotional reaction. Loosely speaking, mine is the view that a person sometimes has *leeway* about just how intensely to take up a fitting emotion.⁹ My degree of admiration for an athlete, for instance, might be fitting, and it might nevertheless be the case that if I had admired that athlete ever-so-slightly more than I actually do, my degree of admiration would still have been fitting. If this is right, then a single evaluative property (here, the athlete's admirability) can make fitting a *range* of fine-grained reactions, all of the same (coarse-grained) type.

Defenders of uniqueness about fitting degrees of emotion, by contrast, sign up to a view on which every fitting emotional reaction must be tailored narrowly to some maximally precise quantum of emotion. This picture is both uncompromising and counterintuitive; the burden of proof lies with its defenders.

Now, there might be ways of meeting that burden of proof. A natural place to look for ideas, here, is the debate over a parallel uniqueness claim:

Uniqueness about Rational Credence Holding fixed an individual's epistemic position, there is at most one rational credence for that individual to bear toward any given proposition.

Defenders of uniqueness about rational credence have marshaled a wide variety of arguments to support it. But, as I'll now demonstrate by considering two examples, those arguments tend to be much less promising as defenses of uniqueness about fitting degrees of emotion than they are as defenses of uniqueness about rational credence.

Perhaps the most prominent line of support for uniqueness about rational credence stems from a worry about *arbitrariness*.¹⁰ Here's one common way to bring out that worry. Suppose, for reductio, that uniqueness about rational credence is false; there are some situations in which a person could be equally rational to take up any of a range of credences toward *p*. If that's right, then whatever credence the person ends up taking up toward *p*, that credence will be, in an important sense, arbitrary. It will be no better, from the perspective of rationality, than certain alternative credences would have been. This might seem destabilizing, especially from the perspective of an agent who acknowledges that her situation is one of this sort. If I learn that any of a range of credences toward *p* would be rational for me, indeed, it might seem that the

⁹ This sort of view is sometimes called "permissivism." But I avoid that label here to avoid the implication that I am analyzing the fittingness of fine-grained emotions using the deontic property of *permission*.

¹⁰ See White (2005, 2013) and Feldman (2007) for defenses of uniqueness motivated in part by worries about arbitrariness.

only way for me to hold one of those credences would be to groundlessly plump for one of the eligible set. But, the thought goes, rational credences cannot be groundlessly selected in this way. So we should discharge our initial supposition; we should accept that uniqueness about rational credence is true.

The arbitrariness worry gains its purchase from a doubt we tend to feel about the rationality of holding a credence when that there's nothing that makes it better than certain alternatives. Ginger Schultheis (2018) offers a distinct argument for uniqueness about rational credence that draws on a different kind of doubt: doubt about the rationality of holding onto a given credence when there *is* something that makes it downright *worse* than other alternatives. Suppose, again for *reductio*, that there is a range of rational credences I could take up toward *p*. Since I'm aware of my own fallibility when it comes to determining the boundaries of this range, Schultheis claims, I should be much more confident that the credences toward the middle of the range are rational than that the credences on the boundaries of the range are rational. But that, says Schultheis, makes it objectionably *risky* for me to settle on a credence on the lower boundary of this range. If I did so, I would be in an uncomfortable position: it would be rational for me to consider other credences much more likely to be rational than my own credence. Schultheis argues that this sort of risk-taking is not just problematic but downright irrational; "it's not rational to adopt a credence that's risky by your own lights when you know of a safer option" (2018, 866). Given this result, she argues, we must reject one of the core stipulations about this case: that there could be a range of rational options for my credences.

Both Schultheis's argument and the arbitrariness worry gain their argumentative power from the unease we're apt to feel when we note that we're holding onto a credence that's flanked by other equally-good—or even better—alternatives. But it's far from clear that we should feel any unease about holding fitting emotions that are flanked by other fitting alternatives.

To bring this point out, consider a metaphor: suppose that I'm trying to finish decorating a room, and I aim to choose curtains that will fit with the room's decorative scheme. I've isolated a range of colors for the curtains—ones that run from baby blue to deep blue—and I hold that any of those colors would fit the decorative scheme of the room. Now, suppose that I plump for a baby blue that's on the edge of the eligible range. In doing so, I might be taking a kind of risk—after all, I should be less certain that this shade of blue matches the room than that the shades in the center of the eligible range do. But it does not follow from the fact that my curtains are a risky choice that they fail to match the room. Nor does it follow that they match the room any less perfectly than would other, "safer" shades of blue.

This is only a metaphor, of course; I do not claim that the sort of "fit" that can obtain between curtains and a room is just the same sort of fit that can obtain between an emotion and its object. I mean, instead, to bring out the following point: while there are some kinds of success (like the success of rationally holding a credence) that do seem to be sensitive to the facts about whether there are other equally-good or better options, there are other kinds of success (like the success that a set of curtains exhibit when they match a room) that are not sensitive to the facts about other options in this

way.¹¹ And the fittingness of a degree of emotion seems like an excellent candidate to be a property of the latter sort, not the former. On a traditional gloss, the fittingness of emotion is not a matter of whether it is one's maximally safe option, or the option that would be recommended by the most reasonable policy. It's simply a matter of whether the emotion matches its object. This seems like the kind of success that a degreed emotion could exhibit even if it does not occupy some special, privileged status relative to a set of eligible options.

I began this section by noting that would-be defenders of uniqueness about degrees of fitting emotion take up a significant burden of proof. Prominent arguments in favor of uniqueness about rational credence seemed like a promising place to look for tools that could help to lift that burden. But, on closer inspection, it's far from clear that any of those arguments can do the trick.¹² So I'll move forward on the assumption that uniqueness about degrees of fitting emotion is false.

2.2 From Unfitting Weak Emotion to Unfitting Apathy

My second thesis is that there are at least some kinds of emotional neutrality—that is, states characterized by the *lack* of emotion—that can be fitting or unfitting in their own right. This second thesis, like my first one, is *prima facie* plausible. What's the fitting response to the fact that the date on which I last clipped my fingernails was a Thursday, or that the time at which I last sneezed was 8:34 AM? Plausibly, a kind of indifference or apathy—the emotional equivalent of a shrug. And what should we say about a person who takes up this very apathy toward an enormously significant event—say, the fact that all their loved ones have been suddenly plunged into serious danger? Here's a very natural thing to say: their apathy is unfitting. Cases like these showcase the common-sense appeal of the claim that emotional neutrality can be fitting or unfitting.

We can motivate the point further by appealing to a distinction, drawn by Justin D'Arms and Daniel Jacobson, between two ways in which an attitude can be unfitting. First, an attitude can have the wrong *shape*; it can be the wrong sort of attitude for the object toward which it is directed (2000, 73). If I find myself feeling irrational anger toward someone who's done nothing wrong (perhaps they simply remind me of someone else who's bullied me), my anger is unfitting in virtue of having the wrong shape. But even when an emotional reaction has the right *shape*, it can still be unfitting in virtue of having the wrong *size* (2000, 74). To say that an emotion has the wrong

¹¹ Berker (forthcoming, sec. 4) offers a similar conclusion, claiming that fittingness (unlike permission and requirement) is “not alternatives dependent.”

¹² There are other arguments for uniqueness about rational credence, too, but they also seem ill-suited to support uniqueness about degrees of fitting emotion. Dogramaci and Horowitz (2016), for instance, argue that uniqueness about rational credence is part of the best explanation for why, when we promote the rationality of our community members, we thereby promote the reliability of testimony. This strategy seems unlikely to translate smoothly to fitting emotion; it's far from clear that there is, or that there could be, a robust social practice of emotional transmission that appropriately parallels the practice of belief-transmission through testimony that is the focus of Dogramaci and Horowitz's discussion.

size for its object is to say that the *degree* of that emotion is unfitting; it is either an overreaction or an underreaction.

Once we note that an emotion can be unfitting in virtue of its size being too small, we have excellent reason to say that at least some forms of emotional neutrality can also be unfitting. That's because there is a continuum that links holding an attitude weakly and (at least some ways of) not holding the attitude at all. We can illustrate that continuum with an example. Start by imagining a character who feels shame quite intensely. Then, imagine weakening that feeling of shame by degrees, until it becomes the weakest degree of shame the person is psychologically capable of feeling. Then, imagine weakening the feeling of shame even further than that. This final weakening of the person's response would eliminate the person's shame altogether. The lack of shame, then, is a weaker reaction than is a halfhearted sort of shame, in just the same way that a halfhearted sort of shame is a weaker reaction than an intense sort of shame. To use a visual metaphor: the absence of shame can be placed at the origin (the "zero point") of an axis that ranks responses by the amount of shame involved.

Now, imagine a case in which a person's shame is *unfittingly* weak. Suppose that, throughout his whole life, Martin Shkreli feels only a tiny twinge of shame for all his misdeeds. The twinge is barely noticeable and lasts for only a moment; he quickly brushes it aside while walking across his apartment to listen to the secret album that he purchased at auction from the Wu-Tang Clan. This reaction would be unfittingly weak; Shkreli's misdeeds were serious, and they merit a more intense reaction. But now imagine a different version of Shkreli, one who considers his past misdeeds in just the same circumstances, but who never reacts with any shame at all. This reaction seems unfitting for just the same reason; Shkreli's misdeeds were serious, and they merit a more intense reaction than he offers. To put the point differently: if it would be unfitting for Shkreli to feel only a tiny twinge of shame, then surely Shkreli cannot avoid the charge of unfittingness simply by feeling no shame at all. Both of those reactions, as we've seen, can be arranged on an axis that ranks responses by the amount of shame involved, and both fall well short of the region on that axis where shame would have the "size" that fits Shkreli's misdeeds.¹³

We don't have to rely only on intuitions about cases, then, to support my second thesis. We can also note that anyone who rejects that thesis—anyone, in other words, who says that emotional neutrality can never be fitting or unfitting—is forced into the awkward conclusion that there is an enormously significant, unbridgeable gap between the normative status of weak emotion and the normative status of emotional neutrality. But the notion of such a gap looks highly implausible. So we should embrace the possibility of fitting and unfitting emotional neutrality.

I've been claiming, throughout this section, that there are at least *some kinds* of emotional neutrality that can be fitting or unfitting. It's worth pausing to note why that qualification is an important one. There are, plausibly, a variety of different properties that could be picked out with the label "emotional neutrality." One is simply the property of lacking emotion. Many people, for instance, have never heard

¹³ [Redacted for blind review]

of the movie *Point Break*, and so do not feel any emotion at all directed toward *Point Break*. There is a sense in which those people are emotionally neutral about *Point Break*, but it is not a very robust sense. After all, rocks and trees are also, in this sense, emotionally neutral toward *Point Break*; they lack any emotions toward it. Call this sort of emotional neutrality *mere emotionlessness*.

It's tempting to think that mere emotionlessness is never fitting or unfitting. (Do we really want to say that a rock has the fitting reaction to the fact that I last sneezed on a Thursday?) But there are other kinds of emotional neutrality as well. There is a difference, for instance, between the person who has never heard of *Point Break* and the person who, upon attentively watching *Point Break*, finds that it leaves her cold. If that's right, then there is a distinction to be drawn between mere emotionlessness and a positive, object-directed state of emotional neutrality. Let's call that positive, object-directed state *apathy*. There are interesting questions about how precisely to characterize apathy, and I'll return to some of those questions in section 3. But for now, I'll use the label to pick out the sort of emotional neutrality that is importantly continuous with weak emotion in the following way: both of them are positively held, object-directed states that can be unfitting in virtue of falling short of the degree of intensity with which it's fitting to hold a particular emotion.

2.3 Bridging the Gap between Weak Emotion and Apathy

My first thesis claims that, in some cases, there is a range of different degrees to which one could fittingly take up a given emotion toward a given object. In defending my second thesis, I noted that the continuum of degrees for emotion shades smoothly, at its lower end, into a kind of emotional neutrality, which I've called *apathy*. My third and final thesis draws on both of these insights. In at least some cases, I claim, the range of fitting fine-grained attitudes toward a given object *bridges the gap* between weak emotion and apathy. This is to say that, in some cases where coarse-grained emotion *E* is a fitting response to a given object, there are a range of different fine-grained reactions that are fitting responses to that object: some including positive degrees of *E* and one (apathy) which includes no degree of *E* all.

To make this proposal more concrete, imagine a garden-variety case where uniqueness about degrees of fitting emotion fails. Suppose, for instance, that I'm listening to gorgeous music, and that I fittingly feel quite a lot of aesthetic pleasure. But also suppose that my reaction would also have been fitting if I'd felt ever-so-slightly less aesthetic pleasure. Anyone friendly to my first thesis, which denies uniqueness about degrees of fitting emotion, should accept that some cases like this exist.

Now consider a second case: one in which I fittingly react to some object with *very weak* aesthetic pleasure. Suppose, for instance, that I observe a somewhat-beautiful cupboard door, and I feel only a tiny twinge of aesthetic pleasure—the smallest iota of aesthetic pleasure that it's psychologically possible for me to feel. Further suppose that, in this case, uniqueness about degrees of fitting emotion is false, for just the reason that it was false in the previous case: namely, I could've felt ever-so-slightly

less aesthetic pleasure than I actually do, and my reaction would still have been fitting. Here's the rub: given the extraordinary weakness of my actual aesthetic reaction, weakening that reaction any further would eliminate my pleasure altogether. It would take me to the zero point on the axis that ranks cupboard-door responses by the amount of aesthetic pleasure involved. In other words, it would amount to my becoming *apathetic* about the cupboard door. If there are any cases of the sort I've just described, then my third thesis is true: at least some of the time, one can fittingly respond to a given object *either* with some degree of emotion *or* with apathy.

Enough illustration; should we think that there *actually are* any cases of this latter sort? In other words, should we think that my third thesis is true? Again, I think that the thesis enjoys quite a bit of initial plausibility. Cases that involve very small misfortunes help to bring out that plausibility. Imagine a case in which you've made a mistake to which you could fittingly respond with the smallest psychologically-possible degree of regret. Perhaps, for instance, you've made a careless banking error, and because of your error, you'll have to pay a \$5 fee. (If that number doesn't seem to you to merit the tiniest possible iota of regret, tweak it until you reach a number that does.) Now, imagine a second case in which you make a mistake that's half as significant. Perhaps you were half as careless, and the banking fee you'll have to pay is half as steep: it's \$2.50. If my third thesis is right, there's room for the following verdict about the revised case: you could fittingly respond to your mistake either with regret or with apathy. But if my third thesis is false, that can't be right: it must be either uniquely fitting to respond to your mistake with a tiny twinge of regret or uniquely fitting to be apathetic.

The former treatment of the case strikes me as a natural and appealing one. Regret, after all, gets something right about the situation. There was a careless error that resulted in a loss, and that's the sort of event that generally merits regret. But apathy also seems to get something right about the situation: loosely speaking, the loss is not a big deal. Each of the two attitudes, then, has a strong claim to matching its object in just the way that's characteristic of a fitting attitude. The burden of proof lies with anyone who wants to rule out the possibility that there could be *any* case of this sort: a case in which a person could fittingly respond to the same object either with emotion or with apathy.

I'll consider one tempting strategy for meeting that burden of proof. An opponent of my third thesis might offer the following argument:

P1. “Fitting” is an absolute gradable adjective.

P2. If “fitting” is an absolute gradable adjective, then any reaction that is less-than-maximally fitting is unfitting.

P3. Whenever an object merits an emotion, apathy is not a maximally fitting reaction to that object.

C. So, whenever an object merits an emotion, apathy toward that object is unfitting.¹⁴

Let’s take these premises in order. The first premise claims that “fitting” is an absolute gradable adjective. This is to say, first, that “fitting” is a *gradable* adjective, which means that it ascribes a property that comes in degrees. Adjectives like “married” and “dead” are non-gradable; no dead person is more dead than another. But adjectives like “expensive” and “empty” are gradable; some vacations are more expensive than others, and some jars are emptier than others.

Gradable adjectives are conventionally divided into two categories: *absolute* gradable adjectives and *relative* gradable adjectives. Absolute gradable adjectives, like “empty,” are associated with *closed* scales; though some jars are emptier than others, there comes a point at which a jar is so empty that no jar could be emptier.¹⁵ Relative gradable adjectives, by contrast, are associated with *open* scales; there’s no point at which a vacation is so expensive that no vacation could be more expensive.

There has been some recent debate as to whether fittingness is gradable. Maguire (2018, 791-2) argues that fittingness does not come in degrees; Berker (forthcoming, sec. 4) argues for the opposite conclusion.¹⁶ But, granting for now the assumption that “fitting” is a gradable adjective, we should acknowledge that it’s better-understood as an absolute gradable adjective than as a relative one. Here’s some evidence: it’s usually infelicitous to modify relative gradable adjectives with modifiers like “perfectly” or “completely.” No vacation is *perfectly* expensive, and no basketball player is *completely* tall. But we don’t find that infelicity with “perfectly fitting” or “completely fitting”; those phrases, like “completely empty,” seem felicitous. That provides some evidence that “fitting” is an absolute, not a relative, gradable adjective.

Let’s move on to premise 2. On a common (but far from uncontroversial) approach to absolute gradable adjectives, any less-than-maximal application of the adjective should be understood as loose, inaccurate usage.¹⁷ On this approach, any jar that is

¹⁴ For a similar argument (about epistemic rationality, not fittingness) see Siscoe (2022).

¹⁵ This exposition elides the distinction between *partial* absolute gradable adjectives like “wet,” which are associated with scales with minimal endpoints, and *total* absolute gradable adjectives like “empty,” which are associated with scales with maximal endpoints. If “fitting” is an absolute gradable adjective, it is of the latter variety.

¹⁶ It’s worth noting that Berker and Maguire seem to agree that the natural-language term ‘fitting’ is gradable; Maguire simply insists that the *property* of fittingness is non-gradable.

¹⁷ Burnett (2017: 70-2) defends this model at length; for philosophical work friendly to this model, see Hawthorne and Logins (2021), Siscoe (2022), and Unger (1979, ch. 2).

less than totally empty is, strictly speaking, not empty; likewise, any gold that is less than perfectly pure is, strictly speaking, impure. If we adopt the same approach to “fitting,” we’ll be inclined to embrace P2, and to say that any reaction that’s less-than-perfectly fitting is, strictly speaking, unfitting.

P3 claims that, whenever an object merits an emotion, apathy is not a maximally fitting response to that object. To see the appeal of this third premise, return to the case of the very small bank fee that I considered above. I suggested that regret and apathy both have a claim to being fitting responses to this bank fee. But a defender of P3 might say that a very weak twinge of regret is *more* fitting than apathy is. After all, very weak regret, unlike apathy, successfully registers the fact that there has been a loss due to a careless error. On top of that, very weak regret also has the virtue that I’ve claimed for apathy; precisely by being a very *weak* form of regret, it registers the fact that the loss is *not a big deal*. Doesn’t this show that a tiny twinge of regret is strictly superior to apathy when it comes to “matching” its object? And doesn’t it therefore also show that regret is *more fitting* than apathy? If so, then apathy is less-than-maximally fitting, and, given premises 1 and 2, it is (strictly speaking) an unfitting response.

The same pattern will hold for any case in which very weak emotion and apathy are both claimed to be fitting responses. Since fitting weak emotion will always reflect something about the world that apathy fails to register, weak emotion will always win out over apathy as *more fitting*, and apathy will turn out not to be a fitting option at all. So there will never be a case in which both emotion and apathy are fitting responses to the same object.

This argument could be resisted in a host of different ways. But its most weakest point is premise 3. The objector assumes that, because fitting weak emotion always reflects *more* about the world than apathy does, fitting weak emotion is always *more fitting than* apathy. But that assumption is a bad one. To bring the point out, I’ll lean on a common metaphor used to illuminate the fittingness of emotion (one which some perceptualists and cognitivists about emotion will treat as literal truth rather than metaphor). According to the metaphor, the fittingness of an emotion is a matter of its representing the world *accurately*, in the same way that a belief, or a photograph, can be said to represent the world accurately.

Here is the point I want to bring out: *more detailed* representation is not necessarily *more accurate* representation. We all agree that, even though some beliefs have contents that reflect more detail about the world than others, that does not make the former beliefs more accurate than the latter ones. And we also all agree that, though the bottom half of a photograph usually represents the world in less detail than does the whole photograph, that does not make the bottom half of the photograph any less accurate than the photograph considered as a whole. We should acknowledge the same point when it comes to our emotional lives: even if we grant that apathy fails to reflect some evaluatively relevant facts about the world, we needn’t conclude that this makes apathy less accurate—or, to discard the metaphor, less *fitting*—than weak emotion would be.

Some might, however, be tempted by that conclusion on other grounds. To return to the metaphor of representation: some might suspect that apathy toward an object that merits an emotion is not just a way of leaving out detail; instead, it's a way of positively misrepresenting the object. But what would the misrepresentation in question amount to? Here's one natural proposal: apathy, in virtue of being the minimally intense reaction one can have to a situation, represents its object as being *maximally evaluatively insignificant*.¹⁸ But there are good reasons to be suspicious of this proposal. There's usually a kind of tension or incoherence involved in holding an attitude while knowing full well that its fittingness-conditions are not met. There's an incoherence, for instance, in feeling resentment toward someone that you know full well has done no wrong, or in feeling proud of something that you know full well doesn't reflect well on you. But there needn't be any incoherence or tension involved in being apathetic about some feature of the world while acknowledging that it is less than maximally evaluatively insignificant. So there are reasons to doubt that apathy is concerned narrowly with the maximally evaluatively insignificant.

In this subsection, I defended my third thesis: there are at least some cases in which a person can fittingly respond to the very same feature of the world either with emotion or with apathy. That claim has quite a bit of initial plausibility, and we haven't found any compelling reasons to doubt it. Let's move on, then, to see how this picture of the relationship between fitting emotion and fitting apathy solves our riddle about the nature of fittingness.

Section 3: Solving the Riddle

On the picture I've defended, there is a unified structure to fittingness. Whenever a coarse-grained emotion is fitting, there are associated facts about the degrees to which it would be fitting to take up that emotion. And at least sometimes, there are a *range* of different degrees to which it would be fitting to do so—in other words, a range of different fine-grained emotions that could be fittingly directed at the very same feature of the world. Further, the degrees to which it's fitting to take up an emotion can in principle be located anywhere on a spectrum from minimally to maximally intense.

Within this unified picture of the structure of fittingness, however, there is room for variation. In some cases, the range of fitting degrees for an emotion stretches far enough to include degree zero—a state of emotionlessness that I've called *apathy*. In those cases, the fittingness of the coarse-grained attitude is permission-like; there are both fitting fine-grained ways of having the emotion and also fitting fine-grained ways of lacking it. This explains why the beauty of a cupboard door sometimes does not demand a response; aesthetic pleasure toward the cupboard door is fitting, but so is apathy toward the cupboard door.

¹⁸ Here's a second one, which faces similar problems: apathy represents its object as being *so insignificant that it couldn't be fittingly responded to with emotion*.

In other cases, however, the range of fitting degrees for a given emotion does not include degree zero. In those cases, the fittingness of the coarse-grained attitude behaves like a demand; the only fitting fine-grained responses all involve, to some degree or other, having the coarse-grained attitude. This explains why the shamefulness of Martin Shkreli's behavior is a demand-like property: any response that does not involve any degree of shame falls outside the range of fitting shame-responses. This means that lacking shame—that is, being apathetic—would be a positively *unfitting* way for Shkreli to respond to his misdeeds.

The model that I've defended, then, is both independently motivated and apt to neatly explain how fittingness could be demand-like in some cases and permission-like in others. But does it fully address the riddle that I raised in section 1? Even readers who sign up to everything I've said so far—agreeing, that is, that apathy can be either fitting or unfitting, and that this marks an important difference between cases of fitting emotion—might suspect that there are at least some example cases regarding which my model offers uncomfortable verdicts.

One such example case is briefly mentioned by Berker (2022). In this example, while you walk past an animal shelter, you notice a lovable cat—that is, a cat that merits, or makes fitting, your love. Berker advocates for a lenient reaction to this case; even if you notice the cat and its lovable features, he suggests, there's no sense in which you go wrong if you do not come to love the cat. My model only offers one way to vindicate lenient reactions of this sort: we'd have to say that the cat merits both emotion and apathy. But that proposal might seem indefensible; is apathy really a good match for the cat's lovable features? Inspired by cases like these, some readers may go looking for different explanations of the gap between cases in which fittingness is permission-like and cases in which fittingness is demand-like.

On closer inspection, however, it's not clear that there are any other viable ways to vindicate a lenient approach to cases like this one. Suppose that someone walks past the cat without coming to love it. Everyone, no matter their view about fittingness, permission, and requirement, must say one of three things about this person's emotional neutrality. First, one could say that it's fitting; this is the option I've just discussed. Second, one could say that it's unfitting; this is also an option highlighted by my model, but it's not a lenient one, and so won't satisfy the objector I'm imagining. Finally, one could say that the person's emotional neutrality is *neither fitting nor unfitting*. This is the only open strategy for my objector. And it's not a promising one.

The first problem with this strategy is that it threatens to overgeneralize. Suppose that there's some form of emotional neutrality—perhaps the *mere emotionlessness* I mentioned in section 2.2—that is simply not apt for appraisal in terms of fittingness. And suppose that we can take up that distinctive kind of neutrality even in cases like the case of Berker's cat; that is, in cases where we're aware of, and we attend carefully to, the fact that a particular object merits an emotional response. Here's the problem: now that this special kind of neutrality is on the table, it's very hard to see why it wouldn't be an eligible way to respond to *any* object that merits an emotional response. If this is all it takes for fittingness to be permission-like, then, we seem pushed toward the conclusion that fittingness is *always* permission-like. So this

proposal misfires. It starts out as an attempt to vindicate the intuition that we should adopt a *distinctive* sort of lenience toward emotional neutrality in certain cases (like the case of Berker's cat), but it ends up doing the opposite: it loses track of the resources necessary to treat any case with greater lenience than it treats any other.

Second, there are good reasons to doubt that we can identify a kind of emotional neutrality with all the features that this proposal requires. I lack the space to argue for this conclusion in detail,¹⁹ but I'll aim to bring out my concern by discussing, in a schematic way, the most promising version of the view that emotional neutrality can be neither-fitting-nor-unfitting.

What could explain why some kinds of emotional neutrality, unlike emotion, lie beyond the pale of appraisal as fitting or unfitting? Many will be tempted to answer in the following way: emotion positively construes or represents its object as having certain evaluative properties. But there is at least some form of emotional neutrality that does not do so.

So far, this view is compatible with everything I've said; I granted, in section 2.2, that one can be *merely emotionless* toward objects that are entirely outside one's ken. Plausibly, this is not a way of construing or representing those objects. But for this proposal to help us offer a novel interpretation of Berker's cat example, it will have to go further. It will have to claim that there are cases in which, even though one is fully aware of an object's emotionally relevant features, one could *still* react with a kind of neutrality that avoids construing or representing the object in the way that emotion would.

On closer inspection, however, it's hard to see how there could be such a kind of emotional neutrality. The only defensible versions of this proposal will have to appeal to a very thin, obscure notion of construal. And it's very hard to say how any kind of emotional neutrality toward an object within one's ken could fail to involve the relevant kind of construal.

To see the problem, start by considering a case in which weak emotion clashes with one's considered evaluative judgments. Suppose, for instance, that I'm faced with a horrifying tragedy while exhausted. I immediately believe, on the basis of a vivid intellectual seeming, that the tragedy merits a great deal of distress. But my exhaustion dulls my emotions, and I only feel a weak twinge of distress. On the proposal we're now considering, an emotion is eligible to be considered unfitting only if it construes its object in a certain light. And, presumably, my weak reaction to the tragedy is *unfitting* precisely because it *misconstrues* its object. Not because it has the wrong shape for the tragedy—its shape is exactly right—but because it has the wrong size. We'll have to say, then, that it construes the tragedy as less distressing than it really is. But in what sense, exactly, does my emotion construe the tragedy as not-very-distressing? Nothing in my mental economy involves a *commitment* to the tragedy as only-slightly-distressing. Nor does it *seem* or *appear* to me, in any sense, that the tragedy is only-slightly-distressing. The only sense in which my emotion

¹⁹ [Redacted for blind review]

construes its object incorrectly, on the face of it, is that it happens to fall short of the fitting level of intense distress.

If this is all that's involved in *misconstruing* the importance of tragedies, however, then emotional neutrality seems apt to misconstrue the importance of tragedies in just the same way. Imagine a variant of the tragic-news case. Suppose that I'm aware of the tragedy in just the same way, and that I remain firmly convinced on the basis of a vivid intellectual seeming that it merits great distress. But suppose that, this time, I happen to be so exhausted that I'm left entirely unmoved. The defender of the strategy we're now considering will have to say that there's at least some version of this case where my emotionless reaction is importantly different from the weak reaction from the original case: unlike the weak reaction, it does not misconstrue the tragedy. But it's hard to see what could justify this claim. The problems with these two reactions may be different, but the difference in question is surely a difference of degree, not a difference of kind.

We set out to find a kind of emotional neutrality with two crucial features: first, it can be taken up even toward objects with which one is fully acquainted, and second, it does not involve the special features that render *weak* emotion evaluative as fitting or unfitting. But, on closer inspection, it seems doubtful that there is any kind of emotional neutrality that fits this description. It takes very little for the degree of one's emotion to be unfittingly weak; the weakness of one's emotion needn't reflect one's judgments, seemings, true self, or indeed any robust person-level commitment to how the world stands.²⁰ This makes it very hard to see what the special ingredient could be that—in cases like the case of the lovable cat—is missing from emotional neutrality, and keeps it from being fitting or unfitting.

The attempt to account for fittingness as a permission-like status by calling neutrality *neither-fitting-nor-unfitting*, then, faces daunting challenges. I doubt that those challenges can be satisfactorily met. Rather than sticking with the attempt to address Berker's lovable cat by making space for a version of emotional neutrality that's *neither-fitting-nor-unfitting*, we should make peace with one of the two options highlighted by my model. In other words, we should acknowledge that the apathy of those who walk past the cat (while noticing its lovable features) is either fitting or unfitting.²¹

I won't insist on either of these options; a defender of my view could safely embrace either. If we insist on a lenient reaction to those cases, we'll have to accept the surprising conclusion that it's fitting to be apathetic toward Berker's cat. But if this seems too disrespectful to the cat, we could reconsider our initial lenient reaction,

²⁰ Clarke and Rawling (forthcoming) argue for a related conclusion about the emotions involved in blame.

²¹ There is an interesting question, and one that deserves further discussion, about just how strong one's epistemic connection to a given object must be before one's neutrality toward that object is eligible to be considered fitting or unfitting. For the purposes of this paper, I'll simply propose the following sufficiency claim: if one is aware of, and carefully attends to, the emotionally significant features of an object, one's neutrality toward that object is evaluative as fitting or unfitting.

and say that the fittingness of loving the cat turns out to be a demand-like property. Those who meet lovable cats and fail to love them are, on this view, falling short: they're unfittingly (albeit blamelessly) apathetic toward an object that does not merit apathy. This is a perfectly acceptable move for those who embrace my model; as I mentioned in section 1, the model I've offered here is not solely motivated by the goal of vindicating intuitive reactions to particular cases. To the contrary, it's also independently motivated by weak and plausible assumptions about the relationship between coarse-grained attitudes, fine-grained attitudes, and apathy.

Conclusion

We have a solution to our riddle; we have an account that nicely explains how fittingness could be demand-like in some cases, but permission-like in others. But we've also gained something of even greater importance than the solution to a riddle: we have a nuanced, flexible, and well-motivated account of the relationship between fitting coarse-grained emotion, fitting fine-grained emotion, and fitting apathy. And we may have learned an important lesson about methodology, too: we can gain a great deal of clarity about the norms that govern our emotional lives by paying closer attention to cases where emotion goes missing.

References

- Berker, Selim. (2002). "The Deontic, the Evaluative, and the Fitting." In C. Howard and R. Rowland (eds.), *Fittingness*. Oxford: Oxford University Press.
- Burnett, Heather. (2017). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford: Oxford University Press.
- Clarke, Randolph and Piers Rawling. Forthcoming. "True Blame." *Australasian Journal of Philosophy*.
- D'Arms, Justin and Daniel Jacobson. (2000). "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research* 61(1): 65-90.
- . (2014). "Wrong Kinds of Reason and the Opacity of Normative Force." *Oxford Studies in Metaethics*, vol. 9, edited by Russ Shafer-Landau, pp. 215-244.
- Dogramaci, Sinan and Sophie Horowitz. (2016). "An Argument for Impermissivism about Evidential Support," *Philosophical Issues*, Vol. 26: pp. 130—47.
- Ewing, Alfred C. (2012). *The Definition of Good*. New York: Routledge. (Original work published 1948.)
- Feldman, Richard. (2007). "Reasonable Religious Disagreements." In Antony, L. (ed.), *Philosophers Without Gods*, pp. 194—214. Oxford: Oxford University Press.
- Gertken, Jan, and Benjamin Kiesewetter. (2017). "The Right and the Wrong Kind of Reasons." *Philosophy Compass* 12/5: 1—14.
- Hawthorne, John and Artūrs Logins. (2021). "Graded Epistemic Justification." *Philosophical Studies* 178, 1845—1858.
- Howard, Christopher. (2018). "Fittingness." *Philosophy Compass* 13 (11).
- Maguire, Barry. (2018). "There Are No Reasons For Affective Attitudes." *Mind* 127(507): 779-805.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391—423.
- Rowland, Richard. (2022). "The Authoritative Normativity of Fitting Attitudes." *Oxford Studies in Metaethics*, vol. 17, edited by Russ Shafer-Landau, pp. 108-137.
- Schultheis, Ginger. (2018). "Living on the Edge: Against Epistemic Permissivism." *Mind* 127 (507): 863-879.
- Siscoe, Wes. (2022). "Rational Supererogation and Epistemic Permissivism." *Philosophical Studies* 179 (2): 571-591.

- Unger, Peter. (1975). *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.
- White, Roger. (2005). 'Epistemic Permissiveness.' *Philosophical Perspectives* 19: 445—59.
- . (2013). "Evidence Cannot Be Permissive." *Contemporary Debates in Epistemology*, 2nd ed. Eds. M. Steup, J. Turri, and E. Sosa. Hoboken: John Wiley & Sons. 312—23.
- Whiting, Daniel. (2021). "Aesthetic Reasons and the Demands They (Do Not) Make." *Philosophical Quarterly* 71 (2): 407—427.