

The Deontic, the Evaluative, and the Fitting

Selim Berker
Harvard University
sberker@fas.harvard.edu

[Draft as of June 23, 2021; to appear, after revisions, in a collection of papers on fittingness on Oxford University Press edited by Chris Howard and Richard Rowland.]

1. Introduction

The normative categories can be split up into several distinct families.¹ The most familiar of these are the deontic and the evaluative categories, or “the right” and “the good” as they are commonly known. The members of each family of categories (i) resemble each other in certain ways and (ii) bear certain relationships to each other that they don’t bear to non-family members. Think here of the ties between *required*, *permitted*, and *forbidden*, on the one hand, and between *good*, *better*, and *best*, on the other. It is these interconnections that make it so natural to split up the deontic and the evaluative into separate families and then ask that well-worn question, “Is the good prior to the right, or the right prior to the good?”

Where do the fittingness categories—*fitting*, *apt*, *warranted*, *admirable*, *fearsome*, *shameful*, and the like—fit into this familiar taxonomy? The overwhelming tendency among contemporary philosophers has been to shoehorn them into the standard bipartite framework. Most authors assume that the fittingness categories are deontic categories, and hence are forms of either requiredness or permissibility (although it is usually left unclear which). Other authors go the other way and assume that the fittingness categories are evaluative categories, and hence are types of goodness and badness, of value and disvalue.

The purpose of this chapter is to argue that both of these tendencies are a mistake. The fittingness categories are neither a subclass of the deontic categories nor a cross section of the evaluative categories. Rather, they constitute a separate family of normative categories, with their own distinctive interrelations, their own distinctive logic, and their own distinctive nature. There are at least three major families of normative categories, and we should resist the urge—tempting though it is—to treat the fitting as simply another variety of the deontic or the evaluative.

¹ As is by now customary, I use ‘normative’ as an umbrella term referring to anything having to do with oughtness, value, virtue, merit, etc. I use ‘category’ as a term of art that encompasses not just properties and relations but also entities, connectives, and anything else that is the worldly analogue of words or concepts, in the way in which the property *being red* is the worldly analogue of the predicate ‘is red’ (a piece of language) and the concept IS RED (a vehicle of thought). I use italics to designate categories, and I occasionally drop the copula ‘being’ when referring to properties and relations, if no confusion results.

2. The Deontic and the Evaluative

I begin with a review of the deontic and the evaluative categories, and of why it is customary to see them as forming distinct families.

The deontic categories include at least the following:

required/obligatory/mandatory/compulsory
permitted/permisible/allowed
forbidden/impermissible/prohibited

ought to/should/must/have to/have a duty to
may/can (when given a deontic reading)

optional/merely permitted (i.e. *permitted but not required*)

The evaluative categories, by contrast, include at least the following:

good
bad
neutral/indifferent

better than
worse than
equal in value to
on a par with (?)

strongly best (i.e. *better than all the rest*)
weakly best (i.e. *at least as good as all the rest*)
strongly worst
weakly worst

at least as good as
no worse than
at least as bad as
no better than

second best
good enough
better than most

Note that both families of categories might come in various “flavors.” For instance, maybe we can distinguish moral, legal, prudential, epistemic, all-things-considered, etc. varieties of the deontic categories. Note, also, that various distinctions can be made within a given category. For instance, maybe we can distinguish final and instrumental badness, attributive and predicative goodness, basic and derivative requirements, directed and undirected duties, etc.

Corresponding to each of these families of categories (properties, relations, entities, etc.) is a family of terms and a family of concepts. So in addition to deontic categories such as *being required*, there are deontic terms such as ‘required’ and deontic concepts such as REQUIRED. And in addition to evaluative categories such as *being good*, there are evaluative terms such as ‘good’ and evaluative concepts such as GOOD. For the most part I shall conduct my inquiry in a metaphysical mode and formulate my claims in terms of categories rather than words or concepts, because I hold that ethics is primarily concerned with investigating the nature of, say, requiredness itself, not with investigating the nature of a certain word in English or of a certain vehicle of thought, and similarly for other normative disciplines. But readers who prefer to approach normative subjects in a semantic (or conceptual) key can feel free to translate most of my claims about categories into

analogous claims about words (or concepts) via semantic (or conceptual) ascent.

There are myriad structural and substantive differences between the deontic and the evaluative categories that justify grouping them together into separate families. I mention here five.²

The first difference: requiredness and permittedness are duals of one another, but no such relation exists among the evaluative categories. This relation of duality takes several forms. On the one hand, requiredness and permittedness are *extensional duals*, in that the following biconditional—or something close to it—holds:

B₁. ϕ -ing is required iff not- ϕ -ing is not permitted.³

But we do not have:

B₁'. x is good iff not- x is not E,

where ‘E’ refers to some other evaluative category. For one thing, we need not hold that the sorts of things that can be good are also the sorts of things that can be negated, so ‘not- x ’ might not make sense for some relevant values of x . And for another, even when ‘not- x ’ does make sense (as it arguably does when x is an action or a state of affairs), that x is good is compatible with not- x being good or bad or neutral: sometimes it is good both to do something and not to do it, and other times it is good to do something without its being actively bad or actively good to refrain from doing it.

Because of their extensional duality, requiredness and permittedness bear certain logical relations to each other that have no analogue for the evaluative categories. But there is another form of duality that leads to an even deeper difference. Requiredness and permittedness are also *definitional duals*, insofar as they are

² See also Smith 2005, §3, and Tappolet 2013. There is one commonly cited difference between the deontic and the evaluative that I do not focus on here, because I think it is a mistake to take this to be genuine difference: I have in mind the claim that deontic categories such as *ought* and *required* are connected to voluntary choice or action in a way that evaluative categories are not. But this alleged distinguishing feature is spurious: there can be things we ought to believe or intend despite belief and intention not being under our direct or even indirect voluntary control (Feldman 2008, Hieronymi 2008, Chuard and Southwood 2009, McHugh 2017).

³ I add the ‘or something close to it’-qualifier to allow that B₁ might not be the best way of formulating extensional duality. One variant formulation adds in a requirement that ϕ -ing be the sort of thing to which norms of the relevant flavor—whether moral, or legal, or whatever—apply:

B₁*. If ϕ -ing is subject to norms of variety x , then (ϕ -ing is required, iff not- ϕ -ing is not permitted).

(An advantage of this formulation: it allows moral error theorists to say both that nothing is morally required and that nothing is morally permitted.) A second variant formulation replaces the appeal to ϕ -ing’s negation with a quantification over ϕ -ing’s alternatives:

B₁**.

ϕ -ing is required iff every alternative to ϕ -ing is not permitted.
(An advantage of this formulation: it allows one to countenance relations of extensional duality even if one holds, say, that the positive state *believing that p* is deontically assessable whereas the negative state *not believing that p* is not.) Analogues of every argument I go on to make for B₁ could also be made with B₁* and B₁**, so I stick with the more familiar formulation in the main text.

interdefinable via their relation of duality, like so:

D₁. ϕ -ing is required \equiv_{df} not- ϕ -ing is not permitted.

D₂. ϕ -ing is permitted \equiv_{df} not- ϕ -ing is not required.⁴

Some authors act as if we must choose between these two definitions: either requiredness is defined in terms of permittedness, or permittedness is defined in terms of requiredness, but we cannot have both. I myself am suspicious of this commonly held thought. There is a deep-rooted symmetry between requiredness and permittedness, and to say that one is more fundamental than the other strikes me as unacceptably arbitrary, in the way in which it is unacceptably arbitrary to hold that possibility is more fundamental than necessity or disjunction more fundamental than conjunction. It is far better, I would say, to hold onto the intuitive thought that these two deontic categories are interdefinable—or, more precisely, are interdefined—and hence to hold that D₁ and D₂ are both true.⁵ And regardless of whether we accept both D₁ and D₂ or only one of them, the crucial point is that no analogous definition of goodness, badness, or any other evaluative category exists.⁶

I have been focusing on the duality of requiredness and permittedness, but this is just one among a number of similar relations that hold among the traditional deontic categories of requiredness, permittedness, and forbiddenness (i.e. impermissibility). Forbiddenness is related to the other two like so:

B₂. ϕ -ing is forbidden iff not- ϕ -ing is required.

B₃. ϕ -ing is forbidden iff ϕ -ing is not permitted.

Moreover, although it is commonly overlooked and lacks a natural prefix-less term in English, there is a fourth central deontic status—let us call it ‘unrequiredness’⁷—that is the dual of forbiddenness and bears corresponding relations to requiredness and permittedness:

⁴ I leave it open whether these definitions are semantic, conceptual, metaphysical, or some combination thereof (maybe each is a semantic definition that entails a corresponding metaphysical definition). For more on metaphysical definition, see Rosen 2015.

⁵ If doing so is in tension with the sorts of foundationalist assumptions usually made by philosophers when they theorize about definitions, then so much the worse for those assumptions, I would say.

⁶ What about strong and weak bestness? Aren’t they evaluative categories that are duals of one another? That is far from clear. To get them to be extensional duals, we need to make contentious assumptions either about the range of ϕ -ing’s alternatives or about how the value of not- ϕ -ing relates to the value of more specific ways of not- ϕ -ing. (The basic problem: how do we establish the right-to-left direction of $\langle \phi$ -ing is strongly best iff not- ϕ -ing is not weakly best \rangle if ϕ -ing has alternatives other than not- ϕ -ing?) Moreover, even if strong and weak bestness turn out to be extensional duals, it is not plausible that they are definitional duals: a thing’s being better than all the rest cannot be defined as its negation’s not being at least as good as all the rest.

⁷ Paul McNamara (2019) calls this fourth deontic status ‘omissibility’, but I have reservations about his terminology that will emerge in due course.

- B₄. ϕ -ing is unrequired iff not- ϕ -ing is not forbidden.
 B₅. ϕ -ing is unrequired iff not- ϕ -ing is permitted.
 B₆. ϕ -ing is unrequired iff ϕ -ing is not required.

B₁–B₆ give us the logical relations that make up the traditional square of opposition. And underwriting these biconditionals are a raft of definitional interconnections beyond D₁ and D₂, such as:

- D₃. ϕ -ing is forbidden =_{df} not- ϕ -ing is required.
 D₄. ϕ -ing is required =_{df} not- ϕ -ing is forbidden.
 D₅. ϕ -ing is forbidden =_{df} ϕ -ing is not permitted.
 D₆. ϕ -ing is permitted =_{df} ϕ -ing is not forbidden.⁸

We find nothing like this structure within the evaluative categories.

A second difference: goodness and badness are non-privative opposites of a certain sort, but no deontic categories are related in this way. Permittedness and forbiddenness are *privative opposites* (or *complements*): each is the absence or lack of the other. To be forbidden is to be non-permitted, and to be permitted is to be non-forbidden. By contrast, goodness and badness are what I call *polar opposites*: each is the inversely charged flipside of the other, not its mere lack or absence. To be bad is to be anti-good, not to be non-good, and to be good is to be anti-bad, not to be non-bad.⁹ No deontic categories are opposites of this sort. Forbiddenness is neither anti-requiredness nor anti-permittedness; rather, it is related to requiredness via D₃ and to permittedness via D₅. Just as duality is a distinctive relation that we find among the deontic categories but not among the evaluative categories, goodness and badness's polar opposition is a distinctive relation that we find among the evaluative categories but not among the deontic categories.

A third difference, which perhaps is the most obvious: goodness and badness are gradable, but the traditional deontic categories are not. The evaluative properties *being good* and *being bad* are gradable, which

⁸ I omit here six additional definitions linking unrequiredness to the other three traditional deontic categories. So in total we have four deontic categories tethered together by six biconditionals and twelve definitions.

⁹ Here the prefix 'anti-' must here be understood as expressing inversion, as in 'antihero' or 'anticlimax', not as expressing adversariality, as in 'antiaircraft' or 'anti-Semite'. Compare Ibram X. Kendi's (2019) notion of an anti-racist, who is not merely someone who opposes racists, but rather is a kind of inverse racist who actively works to level the playing field among the races. Indeed, Kendi's central distinction between an anti-racist and a non-racist is exactly the distinction I mean to be invoking in distinguishing anti-goodness from non-goodness.

involves a package of three related features: first, they have comparative forms (*being better than*, *being worse than*, etc.); second, they have superlative forms (*being best*, *best worst*, etc.); and, third, they can be acted on by *very*, *somewhat*, and other grading modifiers (so there exist the properties *being very good*, *being somewhat bad*, etc.). By contrast, all of the deontic categories we have considered so far are not gradable. It is never the case that one action is “more required” than another, or that one of the available actions is “most permitted,” or that another of the available actions is “somewhat forbidden.”¹⁰

A fourth difference, closely related to the previous one: in addition to the properties *being good* and *being bad*, there is a third, in-between property *being neutral*,¹¹ but no analogous neutral status exists for the deontic categories. The closest candidate here is *being merely permitted* (i.e. *being permitted but not required*), but that isn’t a true neutral, in-between status. Goodness and badness fall on a (perhaps non-scalar, perhaps multidimensional) spectrum with neutrality lying in between, but requiredness and forbiddenness do not fall on a spectrum with mere-permittedness lying in between. If I can save someone’s life by giving them either antidote A or antidote B, we would never say, “Giving that person antidote A is neither required nor forbidden but in between,” in the way in which we might say of a different situation, “That outcome is neither good nor bad but in between.”¹²

Finally, a fifth difference: something’s deontic status depends on what the alternatives are in a way that its goodness or badness doesn’t. Permittedness and requiredness have a property that I shall call *alternatives dependence*: whether ϕ -ing is permitted or required depends, in part, on the normatively relevant properties of

¹⁰ Two comments about gradability. First, my central claim here is about the gradability of categories—of the properties themselves—not about the gradability of words or other linguistic items. Gradability is a metaphysical phenomenon, in addition to being a semantic phenomenon that can be studied by linguists. Second, although philosophers and linguists often implicitly assume otherwise, being gradable is not the same as coming in degrees. Degreedness entails gradability, but gradability does not entail degreedness. If *being F* is gradable but *being more F than* allows cycles, then it will not make sense to reify a thing’s F-ness into degrees. (I owe this point to Angel Navidad.) If *being F* is gradable but *being more F than* is a massively partial ordering, so that there are a large number of cases in which some x is neither more F than, nor less F than, nor equal in F-ness to some y , then it will also not make sense to reify a thing’s F-ness into degrees. And so on.

¹¹ Other common names for this status: ‘being indifferent’ and, especially in the case of attributive goodness, ‘being average’ (as in: “He’s neither a good nor a bad free throw shooter, but merely average”).

¹² That the deontic categories are non-gradable and lack a neutral middle state might seem obvious, but these points are not always appreciated. In a study with over 1,000 citations, psychologists Fiery Cushman, Lianne Young, and Marc Hauser asked their subjects to rate a number of actions “on a scale from 1 to 7, with 1 labeled ‘Forbidden,’ 4 labeled ‘Permissible,’ and 7 labeled ‘Obligatory’ ” (2006, 1083). Even if we set aside the worry that obligatoriness entails permissibility and interpret ‘Permissible’ here as meaning ‘Merely Permissible’, it is difficult to make sense of what Cushman et al. were asking their subjects. What does 3 on their scale signify? A little bit forbidden? What does 5 signify? Permissible but not quite obligatory?

all the alternatives to ϕ -ing.¹³ What matters is not just the case that can be made for ϕ -ing, but also how that case stacks up in comparison to the case that can be made for each alternative to ϕ -ing. So an action that is required in one situation might have the same (non-comparative) properties and yet be forbidden in another, because certain additional courses of action are available. Goodness and badness, by contrast, are not alternatives dependent in this way.¹⁴

Let us take stock. We have uncovered five key differences between the core deontic properties *required*, *permitted*, *forbidden*, and *unrequired* and the core evaluative properties *good*, *neutral*, and *bad*:

	<i>Required/permitted /forbidden/unrequired</i>	<i>Good/neutral/bad</i>
Duality?	Yes	No
Polar opposition?	No	Yes
Gradable?	No	Yes
Neutral state?	No	Yes
Alternatives dependent?	Yes	No

These differences show why requiredness and permittedness are not forms of goodness and goodness not a form of requiredness or permittedness, and similarly why forbiddenness and unrequiredness are not forms of badness and badness not a form of forbiddenness or unrequiredness. Thus we can rule out the most straightforward way of viewing deontic properties as evaluative properties, or vice versa. But we have left open the possibility of other, more complex proposals. For instance, what if we take requiredness to be a form of bestness? Would such a proposal show that the deontic categories and the evaluative categories are not separate families after all?

¹³ Johann Frick's (MS, 8) term for much the same thesis is "comparativism" (about a given category).

¹⁴ It might seem that at least some first-order moral theories deny that an action's deontic status is necessarily an alternatives-dependent matter; think here of the "absolutist" view on which intentionally killing another innocent person is always impermissible, even if every available alternative would lead to billions of deaths. Personally I find absolutism so outrageously implausible that I am happy to set it aside. But for those who do not wish to rule it out, a version of my fifth difference can survive. Even if something's deontic status is not *always* an alternatives-dependent matter, it *typically* is, but the same is not true of a thing's goodness.

A different sort of objection: if to be good is to be better than most of the items in some contextually salient comparison class (as some linguists propose), does this mean that goodness is alternatives dependent after all? No, for that comparison class need not consist in the alternatives to the thing being evaluatively assessed. Moreover, even if *being good* does turn out to be alternatives dependent for this reason, a version of my fifth difference remains, for *being better than* would still be an alternatives-independent matter: that x is better than y does not depend on the normatively relevant properties of the alternatives to x and y .

No, it wouldn't. To see why, let us consider one salient group of views on which it might be thought that requiredness turns out to be a kind of bestness, namely *maximizing forms of consequentialism*. It follows on such views that requiredness is co-extensive with some manner of (strong) bestness, so that something like the following biconditional holds:

C₁. ϕ -ing is required iff ϕ -ing leads to the (strongly) best outcomes.¹⁵

But this biconditional on its own is not enough to establish that requiredness is a form of bestness, for biconditionals are cheap. Instead, we need something stronger. What might that stronger thesis be? I have argued elsewhere (Berker 2018, 2019) that we do best to see maximizing consequentialists as committed to

C₂. Necessarily, ϕ -ing is required iff, and because, ϕ -ing leads to the best outcomes.

However, C₂ is not a view on which requiredness turns out to be an evaluative category after all; rather, it is a view on which requiredness turns out *to depend on* an evaluative category, and dependence is irreflexive. So on the most natural way of understanding maximizing consequentialism, it does not lead to a collapse of the deontic into the evaluative.

What, though, about other ways of interpreting maximizing consequentialism? Suppose that instead of C₂ we have one or more of the following:

C₃. The property *being required* = the property *leading to the best outcomes*.

C₄. The concept IS REQUIRED = the concept LEADS TO THE BEST OUTCOMES.

C₅. The phrase 'is required' means 'leads to the best outcomes'.

Now to start with, I think these are all implausibly strong ways of interpreting the consequentialist's commitments (Berker 2018, §3). But even if we were to accept one or more of these theses, I don't think we would stop distinguishing deontic categories (or concepts, or terms) from evaluative categories (or concepts, or terms): we would simply come to see the former as a particularly interesting subset of the latter. The deontic and the evaluative would remain distinct in one sense—they are different families—while not being distinct in another sense—they overlap with one another.¹⁶

¹⁵ The details of this biconditional will of course vary depending on the variety of maximizing consequentialism at issue. (See Berker 2013, §2, for a survey of the possibilities; there I use the term 'teleology' where I now prefer 'consequentialism'.) Henceforth by 'best' I mean 'strongly best' unless I specify otherwise.

¹⁶ A similar result holds if we accept

3. The Fitting

Let us turn now to our main topic: the fitting. The fittingness categories include at least the following:

the thin (or basic) fittingness categories:

fitting/ appropriate/ apt/ warranted/ justified/ merited/ deserved/ worthy/ suitable/ proper/ correct (?)
/ called for/ to be done (or *felt, or held*)
unfitting/ inappropriate/ inapt/ unwarranted/ unjustified/ unmerited/ undeserved/ unworthy/ unsuitable
/ improper/ incorrect (?) */ uncalled for/ not to be done* (or *felt, or held*)

the thick (or derivative) fittingness categories:

(un)acceptable/ admirable/ (dis)agreeable/ contemptible/ credible/ deplorable/ (un)desirable/ (un)enviable
/ (in)excusable/ lamentable/ (im)plausible/ preferable/ (un)reliable/ (ir)responsible/ etc.
awful/ disgraceful/ dreadful/ shameful/ useful/ etc.
(un)amusing/ annoying/ disgusting/ (un)interesting/ (un)surprising/ etc.
(un)attractive/ (in)offensive/ (un)persuasive/ repulsive/ suggestive/ etc.
bothersome/ fearsome/ gruesome/ irksome/ tiresome/ etc.
blameworthy/ choiceworthy/ noteworthy/ praiseworthy/ (un)trustworthy/ etc.
(un)funny/ scary/ tasty/ etc.

As with the deontic and evaluative categories, each of these categories might come in different “flavors”: moral appropriateness, epistemic warrantedness, and so on. Some authors use ‘fittingness’ only to refer to varieties of the above normative categories that hold in virtue of the internal or constitutive standards of the item being assessed. That is not how I am using the term. An inappropriate comment need not be inappropriate in virtue of violating “the constitutive standards of comments” in order for its inappropriateness to qualify as a fittingness category—to qualify as being concerned with how that comment, as it were, fits the situation. This metaphor of fit, of something matching or suiting a given circumstance or object, links all of these categories together and makes it natural to view them as a unified group that it is handy to label with the term ‘fitting’.

I have split up these categories into two groups, the thin (or basic) ones and the thick (or derivative)

C₆. It is part of the essence of *being required* that something which is required leads to the best outcomes together with

E. If category *x* features an evaluative category in any part of its essence, then *x* is itself an evaluative category.

But, once again, I think this possibility need not overly concern us, since neither C₆ nor E strike me as particularly plausible. (Concerning C₆: I do not see why our consequentialist thesis is part of the essence of *being required* on its own, rather than part of the collective essence of *being required* and *being best*, taken together, or not an essential truth at all. Concerning E: that a category features a logical [or natural] category within its essence does not thereby make that category a logical [or natural] one, so why should things be any different with evaluativeness? Why does the taint of the evaluative spread through essences, but the taint of the logical [or the natural] does not?) And whether a similar result holds if we have

C₇. ϕ -ing is required \equiv_{df} ϕ -ing leads to the best outcomes

will depend on whether we are working with a conception of definition on which it entails C₂ or instead entails one of C₃–C₅.

ones.¹⁷ The properties *apt*, *appropriate*, *merited*, *deserved*, *called for*, etc. are all, broadly speaking, forms of thin fittingness, in the way that *obligatory*, *mandatory*, *ought*, *should*, *must*, etc. are all, broadly speaking, forms of requiredness. Some of these may even be the same category picked out in different ways: it is not clear there is a difference between being appropriate and being apt, in the way that it is not clear there is a difference between being required and being mandatory. But in other cases, there obviously are differences. For instance, to say that something is warranted is not the same as saying that it is deserved; believing some proposition might be warranted, given one's evidence, but no proposition deserves to be believed (Howard 2018, 7). However, a similar thing is also true in the case of the deontic categories: to say that one ought to ϕ is not the same as saying that it is obligatory for one to ϕ . But this difference between *ought* and *obligatory* is no bar to classifying both as forms of requiredness—to classifying both as boxes, not diamonds, in the familiar notation. Similarly, the fact that *being warranted* and *being deserved* are not mutually entailing is no bar to classifying both as forms of thin fittingness. 'Warranted' and 'deserved' can both be thin fittingness terms without being synonyms.

The thick fittingness properties are analyzable in terms of the thin ones, as was pointed out by Richard Brandt (1946) three quarters of a century ago. Brandt only explicitly applies his analysis to fittingness properties picked out by adjectives ending in '-able'/'-ible', but his analysis can be straightforwardly extended to all thick fittingness properties, like so:

x is admirable =_{df} it is fitting to admire x .

x is contemptible =_{df} it is fitting to have contempt for x .

x is desirable =_{df} it is fitting to desire x .

x is reliable =_{df} it is fitting to rely upon x .¹⁸

x is annoying =_{df} it is fitting to be annoyed by x .

¹⁷ My primary labels for these two groups are 'thin' and 'thick', but I use those terms with some trepidation, because it is controversial whether ADMIRABLE, SHAMEFUL, and the like are thick concepts in the sense made popular by Bernard Williams (1985). Some authors deem them to be (Anderson 1993, 98; Tappolet 2004; Kyle 2020), others deem them not to be (Gibbard 1992; Suikkanen 2009, 778, n. 20), and yet others express uncertainty (Väyrynen 2021). I am happy to switch to a different pair of labels if the doubters are right.

¹⁸ Note that this analysis shows that 'reliable' is a normative term that is part of the fittingness family. So when epistemologists treat 'reliable' as a non-normative term meaning 'truth-conducive', they are using 'reliable' as a technical term, not using it in its ordinary sense. (In everyday speech, to say that Boston's subway system is unreliable is to normatively assess that subway system, not to say something about whether it promotes true belief.)

x is attractive \equiv_{df} it is fitting to be attracted to x .

x is fearsome \equiv_{df} it is fitting to be afraid of x .

x is trustworthy \equiv_{df} it is fitting to trust x .

I have formulated these analyses using the word ‘fitting’, but I just as easily could have used a number of other thin fittingness terms instead: for something to be admirable is for it to merit admiration, or for it to be worthy of admiration, or for admiring it to be apt, or for it to call out for admiration, or Often our term in English for a given thick fittingness property is derived via suffixation from a noun or verb in English expressing the reaction whose fittingness is at issue, but not always. Sometimes the process of suffixation occurred in another language without the root being passed along to English as well, as in the case of ‘risible’ (‘that which merits laughter’, from the Latin ‘ridere’ meaning ‘to laugh at’) or ‘culpable’ (‘that which deserves blame’, from the Latin ‘culpare’ meaning ‘to blame’). And in a few cases, we have no ready way to characterize the relevant reaction except via the thick fitting property being analyzed. For instance, although ‘plausible’ originally referred to that which it is fitting to applaud, that meaning is obsolete, and ‘plausible’ now refers to that which it is fitting to . . . well, find plausible.

In an important and influential series of articles, Justin D’Arms and Daniel Jacobson call those who advocate analyses of the above sort “neo-sentimentalists” (D’Arms & Jacobson 2000a, 2000b, 2006a, 2006b, 2017; D’Arms 2005), but in my opinion that is a regrettable piece of terminology which should be avoided. Brandt-style analyses of thick fittingness properties are entirely independent of the sentimentalist program in metaethics, for two reasons. First, a similar analysis holds for various thick fittingness properties that have nothing to do with the sentiments, such as:

x is punishable \equiv_{df} it is fitting to be punished for x .

x is hireable \equiv_{df} it is fitting to hire x .

x is credible \equiv_{df} it is fitting to believe x .

x is useful \equiv_{df} it is fitting to use x (for such-and-such purpose).

x is persuasive \equiv_{df} it is fitting to be persuaded by x .

x is choiceworthy \equiv_{df} it is fitting to choose x .

x is seaworthy \equiv_{df} it is fitting to sail x .

Hence definitions of this sort are not distinctively sentimental in any way. Second, even those Brandt-style analyses that do appeal to sentiments are compatible with facts about the fittingness of sentiments being fully grounded in facts having nothing to do with the sentiments, such as facts about the promotion of pleasure and pain or facts about the import of the actions those sentiments dispose us to perform. But a metaethical view whose appeal to sentiments is itself explained in terms of something deeper is not one that is aptly described as sentimental, I would say.¹⁹

Not all ‘-ible’/‘-able’ adjectives in English are normative terms that pick out thick fittingness properties. Some are descriptive terms that pick out potentiality properties, such as:

x is visible \equiv_{df} x is able to be seen.

x is detachable \equiv_{df} x is able to be detached.

And some ‘-ible’/‘-able’ adjectives can be used to denote either a thick fittingness property or a descriptive potentiality property. “That sentence is unprintable” can mean “That sentence is not fit to be printed” (because, say, it is too offensive), and it can mean “That sentence is not able to be printed” (because, say, it is too long to write down using existing technology).²⁰

Interestingly, negative prefixes work differently for some thick fittingness terms than others. Often the negative prefix acts as a wide-scope negation:

x is unacceptable \equiv_{df} it is not the case that (it is fitting to accept x).

Or perhaps instead:

x is unacceptable \equiv_{df} x is the sort of thing that can be accepted, and it is not the case that (it is fitting to accept x).²¹

But sometimes the negative prefix acts as a narrow-scope polarity-flipper:

¹⁹ D’Arms and Jacobson also call properties amenable to Brandt-style analyses ‘response-dependent properties’, and others have followed them in this practice (Tappolet 2013, Shoemaker 2017), but in my opinion that is another regrettable piece of terminology which should also be avoided. To say that something is *dependent on which of our responses are merited* is very different from saying that it is *dependent on our responses*, and ‘response-dependent’ only makes sense as a label of the latter, not the former. No one would ever say that punishable offenses have a ‘punishment-dependent’ property because they merit punishment.

²⁰ See Kjellmer 1986. There is also a third class of meanings for ‘-ible’/‘-able’ adjectives: a small number of them refer to a certain ready tendency. For example, ‘knowledgeable’ means ‘tends to readily know things (about the domain in question)’, and ‘perishable’ means ‘tends to perish quickly’.

²¹ A reason to prefer this analysis over the former one: then we need not say that rocks and other things that cannot be accepted are unacceptable. Alternatively, *that x is the sort of thing that can be accepted* might be an enabling condition for or presupposition of unacceptability, without being part of its definition. But regardless of how we implement this thought, that some such qualifier is needed seems to be a general feature of negative prefixes that act as wide-scope negations, regardless of whether they are applied to normative or non-normative terms. Even though rocks are not afraid, it doesn’t follow that rocks are unafraid.

x is undesirable \equiv_{df} it is fitting to be adverse to (i.e. to anti-desire) x .

Not all non-desirable outcomes are undesirable; what makes an outcome undesirable is that it positively calls out for a negative orrectic attitude, not that it fails to call out for a positive one.

A large percentage of the normative terms that we use on a daily basis are thick fittingness terms formed via suffixes like ‘-able’, ‘-ible’, ‘-ing’, ‘-ive’, etc. and their analogues in other languages: these terms are *everywhere*, which makes it all the more surprising how under-theorized they are. Moreover, thick fittingness terms and the properties they pick out are all over the place in philosophy: we regularly appeal to them in our theories (some examples: *accountable*, *admissible*, *answerable*, *attributable*, *culpable*, *eligible*, *excusable*, *liable*, *reasonable*, *reliable*, *choiceworthy*), and we constantly use them to evaluative each other’s arguments and proposals (some examples: *[im]plausible*, *[un]compelling*, *[un]convincing*, *fascinating*, *interesting*, *[un]promising*, *surprising*, *impressive*, *[un]persuasive*, *suggestive*). But since thick fitting properties can be analyzed in terms of thin fittingness properties, I will primarily be focusing on the thin ones in what follows.

4. Fittingness Is Not Deontic

Here is a common thought: the thin fittingness properties such as *fitting*, *apt*, *warranted*, *merited*, and the like are themselves deontic properties, and in particular are a type of permittedness or requiredness. This thought is probably the most widespread assumption made about the typology of fittingness categories. For instance, it is simply baked into the wrong-kind-of-reason literature, where it is standardly assumed that the label ‘fitting-attitude theory of value’ applies equally well to views that analyze (or explain, or otherwise account for) goodness in terms of what one *ought* to value or desire and to views that analyze (or explain, otherwise or account for) goodness in terms of what it is *fitting* to value or desire.²² I myself used to hold this common thought, but I now believe it to be mistaken. I present here five reasons to deny it.

First of all, fittingness and requiredness have different guiding metaphors. (This is a warm-up reason that on its own is not probative.) The guiding metaphors behind requiredness or oughtness—which also happen to be linked to these terms’ etymological origins—concern a person being bound, tied, or forced to

²² See, e.g., Rabinowicz and Rønnow-Rasmussen 2004, Danielsson and Olson 2007, Bykvist 2009, Suikkanen 2009, Jacobson 2011, Tappolet 2013, and Gertken and Kiesewetter 2016.

do something. When we are required to do something, it is as if there is a magic bond to that thing which compels us to do it, and when we are permitted not to do that thing, no such bond is present. Compulsion, bondage, being forced: these are all natural metaphors to reach for when discussing requirements. But not so in the case of fittingness; here the guiding metaphors are quite different. One natural metaphor for fittingness: a puzzle piece locking into place. Another natural metaphor: a key fitting into a lock. These are powerful images that are easily associated with all of the standard locutions used to pick out fittingness properties. When a certain emotional reaction is merited, it is the edge piece that fits the current situation. When some action is called for, it is the key that goes into the current lock. And so on. But neither of these metaphors—the puzzle piece and the key—are natural ones to invoke when discussing duties, requirements, permissions, and other uncontroversially deontic categories. Why would that be so, if fittingness were simply a form of permittedness or requiredness?

A second, more substantive reason to deny our common thought is this. When the time comes to decide whether fittingness is a form of permittedness or instead a form of requiredness, we find that neither proposal quite works: in some cases, taking fittingness to be a type of permittedness seems too weak a proposal, and in other cases, taking fittingness to be a type of requiredness seems too strong.

When I do something shameful, shame on my part is fitting or called for. When you hear a persuasive argument, it is fitting or appropriate for you to be persuaded by it. When a trustworthy person tells you something, it is fitting or apt for you to trust what they say. In none of these cases is it plausible to interpret the normative category at issue as being a form of permission. It is far too weak to say that we are simply *permitted* to be ashamed at our shameful conduct, to be persuaded by persuasive arguments, or to trust what trustworthy people tell us. To say that shame is fitting is not merely to say that being ashamed is allowed—is not forbidden—but rather to say something more.

So maybe think fittingness is instead a form of requiredness. But then we run into the opposite problem in numerous other cases. When I go to the rescue center and am confronted with thirty lovable cats waiting to be adopted, it is far too strong to say that I am *required* to love each cat. Similarly, as we go about our lives, we are not required to laugh at every risible situation, to use every useful tool, to sail every seaworthy vessel, to believe every credible proposition, to hire every hireable candidate, or to be appalled by every appalling

politician. Now it might seem that I am raising a demandingness concern here, but that is not the central issue. Although cases in which it would be difficult or even impossible to have all the fitting responses open to us at a given time make vivid the worry I mean to be raising, that worry is not one that would disappear if we had unlimited cognitive, emotional, and physical resources. Even if there is only one lovable cat at the rescue center, it is still implausible to say that I am required to love it. When someone tells a cringeworthy joke, it is not compulsory for me to cringe, but that is not because cringing would expend crucial emotional resources that could be devoted elsewhere. If I happen upon the last seaworthy ship left in the world, I am under no obligation to sail it, even if I am bored out of my mind and have nothing better to do with my time. As a first-order normative matter, the idea that fittingness takes the form of a requirement is highly dubious.

My argument here has proceeded by (1) assuming a Brandt-style analysis on which something is shameful when it is fitting to be ashamed of it, lovable when it is fitting to love it, cringeworthy when it is fitting to cringe at it, and so on; (2) assuming that fittingness, if it is indeed a deontic category, must take the same deontic shape in each instance of this analysis; and (3) finding some cases where taking that shape to be a form of permission seems too weak and other cases where taking that shape to be a form of requirement seems too strong. A natural strategy for resisting this argument would be to give up on providing a uniform analysis of the thick fittingness properties, in one of two ways. The first is to reject (1) by replacing Brandt's analysis with a mixed analysis on which some thick fittingness properties are given a Brandt-style analysis but other thick fittingness properties are given an analysis that is the dual of his, so that we have:

x is useful \Rightarrow_{df} it is fitting to use x (for such-and-such purpose).

x is shameful \Rightarrow_{df} it is not fitting to not be ashamed of x .

x is attractive \Rightarrow_{df} it is fitting to be attracted to x .

x is persuasive \Rightarrow_{df} it is not fitting to not be persuaded by x .

x is cringeworthy \Rightarrow_{df} it is fitting to cringe at x .

x is trustworthy \Rightarrow_{df} it is not fitting to not trust x .

The second way is to hold onto Brandt's analysis across the board but to reject (2) by offering a mixed account of the form of thin fittingness that gets slotted into that analysis, so that for some thick fittingness properties it is a type of permission—call it 'p-fittingness'—and for other thick fittingness properties it is a type of

requirement—call it ‘r-fittingness’. Then we would have:

x is useful $=_{df}$ it is p-fitting to use x (for such-and-such purpose).

x is shameful $=_{df}$ it is r-fitting to be ashamed of x .

And so on for the others.

Both of these mixed proposals are deeply implausible. There is no evidence that the English suffixes ‘-ful’, ‘-ive’, ‘-worthy’, and the like function in these two importantly different ways.²³ It would be one thing if one of these suffixes lent itself to a permissive analysis and another lent itself to a requiring analysis, or if one thin fittingness term (‘fitting’, perhaps) were more naturally construed as picking out a form of permittedness whereas another thin fittingness term (‘merited’, maybe) were more naturally construed as picking out a form of requiredness. But that is not what we find, so the current proposal requires us to construe the relevant definitions to be metaphysical definitions about the nature of the properties themselves and then to posit that the normative language we have right now is severely out of lockstep with normative reality—always a dangerous claim to make, especially if our language resists adjusting itself over time so as to take on board this alleged distinction in the way things are (as I suspect it will). The first proposal—on which some thick fittingness properties are given an analysis that is the dual of Brandt’s—also faces the problem that such an analysis, with its two extra negations, is implausible on its face. Intuitively, persuasiveness is about the fittingness or aptness of being persuaded by the object in question, not about the aptness of failing to be persuaded by it, and shamefulness is about whether shame itself is merited or called for, not about the normative status of abstaining from shame.

Even worse, these mixed proposals do not succeed in addressing the very problem they were designed to address, for many of the cases we have been considering in which the requiring reading is inadequate are *also* ones in which the permissive reading is problematic. Let us go back to one of those lovable cats in the rescue center. Yes, it is too strong to say that I am required to love that cat. But it is also too weak to say that the cat’s lovability simply consists in my being permitted to love it and nothing more. In calling this animal

²³ Nor is there evidence that ‘fitting’ is polysemous and has both a permissive and a requiring reading. Indeed, the evidence very much points in the other direction. If ‘fitting’ were polysemous in the specified manner, then “It is fitting for me to be ashamed and for you to be angry” should sound zeugmatic; but it doesn’t. So ‘fitting’ fails the conjunction-reduction test for polysemy. It also fails many other tests standardly used to detect polysemy or ambiguity, such as the ellipsis test and the contradiction test. (For more on these tests, see Sennett 2016.)

lovable, I am not just saying that it has qualities that make loving it permissible—that make loving it not forbidden, in the way in which I am not forbidden from tying my right shoe before my left one. Rather, I am saying something more about how the cat’s qualities positively call out for or merit a certain affective response on my part. (Here I am falling back on using other fittingness locutions to describe the relation I am after that is neither a permission nor a requirement, but I suspect there is no hope of doing otherwise.²⁴) Reflection on this sort of case leads me to conclude that construing fittingness as a form of permittedness is *always* too weak. The appallingness of appalling politicians is not exhausted by our being permitted to be appalled at them and their actions; rather, there is a further way in which our being appalled suits or fits what they have done. Useful tools are not simply ones we are allowed to use for the purposes for which they were built. Praiseworthy conduct is worthy of praise, but the fittingness of this praise consists in more than such praise being not prohibited.

These problems with mixed proposals and with uniformly permissive proposals might lead us to go back and revisit the prospects for a uniformly requiring proposal. Maybe the solution is to hold onto Brandt’s analysis of thick fittingness properties across the board and also to hold that fittingness is indeed a form of requiredness, but to add more structure to the way in which we connect fittingness to requiredness, like so:

ϕ -ing is fitting =_{df} if one considers whether to ϕ , one is required to ϕ .

The basic idea is that while we may not be required to believe every single credible proposition out there, perhaps we are required to believe every credible proposition that we consider. Similarly, although we might not be required to be appalled by every appalling thing that someone has done, perhaps we are required to be appalled at an appalling action whenever we attend to it and its appallingness. However, there are at least three problems with this proposal. First, it doesn’t help with many of our original cases where construing fittingness as a permission seems too weak. So it’s fine for me not to be ashamed of my shameful conduct, provided I never consider whether to be ashamed of it? Second, although this proposal does help with some of the cases where construing fittingness as a requirement seems too strong, it doesn’t help with all of them. I am not required to sail every seaworthy ship I consider sailing, to hire every hireable candidate I consider

²⁴ Compare Scanlon (1998, 17) on explaining the reason relation in terms of the counts-in-favor-of relation and vice versa.

hiring, to love every lovable pet I consider loving. Third, if this is all fittingness comes to, then the normative pressure to ϕ that accrues when ϕ -ing is fitting completely disappears the instant one stops considering whether to ϕ . But surely something is going wrong if, for no reason at all, I stop admiring an admirable person a millisecond after I cease considering the question of whether to admire that person. Similarly, suppose I drop my belief in some credible proposition a split second after I bring my deliberations about whether to believe it to a close. It is simply not plausible that fittingness is a form of normative pressure that gets exerted only while one is considering some question. And it is also not plausible to revise the antecedent of our proposed analysis so that it is in the past rather than present tense, like so:

ϕ -ing is fitting \equiv_{df} if one has considered whether to ϕ , one is required to ϕ .

My having once considered some question, several decades ago, does not make me permanently required to respond in a certain way to some fitting thing. (Am I really required to right now believe every credible proposition I considered at some time or other while a teenager?) So perhaps we should instead say that the normative pressure that arises due to consideration of a question dissipates after a certain amount of time. But what could possibly determine that length of time? Does the normative statute of limitations expire after several minutes, several days, several years? In short, taking fittingness to be a requirement that kicks in only during deliberation makes it too ephemeral, taking fittingness to be a requirement that kicks in after any bout of deliberation makes it too permanent, and taking fittingness to be a requirement that kicks in during deliberation and hangs around for some specified amount of time forces us to make arbitrary distinctions.

Thus we are left with a problem: neither permittedness nor requiredness is the right, as it were, fit for the fittingness categories. Does this mean that fittingness is an in-between category that is logically stronger than permittedness and logically weaker than requiredness? I don't think that is the right conclusion to draw. Compare the relation between goodness and permittedness/requiredness. It would be a mistake to hold that *being good* is logically stronger than *being permitted* and logically weaker than *being required*. For one thing, some good things are not even the type of entity that can be permitted or required, so *being good* does not entail *being permitted*. For another, there can be required things that fail to be good, as might happen when some obligatory action is the best of a bad lot, so *being required* does not entail *being good*. Rather than holding goodness to be a logically intermediary between requiredness and permittedness, we are better off holding goodness to be an

altogether different normative category from requiredness and permittedness that doesn't bear any straightforward logical relations to either. That, I think, is the proper way to think of the relation between the thin fittingness categories and the deontic categories. It is not a logical truth that fitting responses are the sorts of things that have a deontic status such as permitted or required, and it is not a logical truth that being fitting is entailed by being required and entails being permitted.

Let us turn, then, to a third, equally powerful reason to deny that fittingness is a deontic category: whereas permittedness and requiredness have duals (namely, each other), fittingness does not have a dual. We have no word in English for a normative category that is the dual of *fitting*, *apt*, *appropriate*, *warranted*, *merited*, *correct*, and the like. *Unfitting* is not the dual of *fitting*: it is, rather, *fitting*'s complement. Similarly, *incorrect* stands to *correct* as *impermissible* stands to *permissible*, not as *required* stands to *permissible*.

Perhaps, it might be replied, *fitting* does have a dual, and it is simply an accident that English lacks a handy term to pick it out. Moreover, the reply continues, we can rectify this situation by introducing a new expression, 'schmitting to ϕ ', meaning 'not fitting to not ϕ .' But this is a cheap way of trying to secure duality. By a similar move, we could argue that *any* property has a dual. For instance, we can no longer say that a crucial difference between *required* and *good* is that one has a dual whereas the other doesn't, because we could always introduce a new expression that supposedly is the latter's dual. But more importantly, this trick doesn't even work. Adding 'schmitting' into our lexicon does secure extensional duality by brute force: it is schmitting to ϕ iff it is not fitting to not ϕ . But the proposed linguistic innovation fails to get us definitional duality. Although we might perhaps have:

It is schmitting to $\phi =_{df}$ it is not fitting to not ϕ ,

it is not at all plausible that we have:

It is fitting to $\phi =_{df}$ it is not schmitting to not ϕ .²⁵

This situation stands in contrast to the situation with permittedness and requiredness, where each definition

²⁵ The implausibility of this definition is clearest if we take it to be semantic or conceptual in nature. Could it really be that one cannot know the meaning of 'fitting' or possess the concept FITTING without also possessing the concept SCHMITTING, in the way in which one cannot know the meaning of 'bachelor' or possess the concept BACHELOR without also possessing the concept MARRIED? However, the needed definition is nearly as implausible if we take it to be a metaphysical definition. Surely there was not a crucial part of fittingness's nature that you were unable to articulate until you were introduced to this new term, 'schmitting'.

of the other via an application of duality is on the same footing. In short, permittedness and requiredness are interdefinable, but fittingness and schmittingness are not. So when it comes to the deepest form of duality, a key difference remains even after we have introduced our new term, for fittingness—unlike permittedness and requiredness—has no definitional dual, and creating new words cannot change that fact.

A fourth reason to hold that fittingness is neither a form of permittedness nor a form of requiredness is that whereas those two deontic statuses are non-gradable, it is at least arguable that fittingness is gradable. The thick fittingness properties—*desirable*, *interesting*, *attractive*, *tiresome*, *trustworthy*, and the like—are all undeniably gradable. One outcome can be less desirable than another, some point can be the most interesting one made on a given occasion, one person can be somewhat attractive while another is very attractive, and so on. It is a delicate question, though, whether the thin fittingness properties are gradable. Some seem not to be. For example, consider *correct*. When Alan Dershowitz defended his apparently shifting views about the criteria for presidential impeachment by saying, “I am . . . far more correct now than I was then,” what he said was not just legally questionable but also grammatically suspect.²⁶ Other thin fittingness properties are commonly taken to be gradable. For example, consider *justified*. Epistemologists often tell us that epistemic justifiedness comes in degrees, and hence is gradable. Finally, a number of thin fittingness properties are borderline cases. For example, consider *appropriate*. Is okay to say, “A tuxedo would be more appropriate for this event than a coat and tie”? I myself go back and forth in my assessment.

However, if the thick fittingness properties are gradable, doesn't that settle the issue of whether the thin fittingness properties are gradable? How could the thick ones be gradable while the thin ones are not? Barry Maguire (2018) has recently proposed a way of taking the thick fittingness properties to be gradable while holding that the thin ones are not. His suggestion is that we analyze, for example, desirability in a standard Brandt-style manner but then analyze greater desirability in terms of the fittingness of a greater level of desire, rather than in terms of some desire's greater level of fittingness, so that we have:

x is desirable $=_{df}$ it is fitting to desire x .

x is more desirable than y $=_{df}$ it is fitting to desire x more than one desires y .

²⁶ <https://www.cnn.com/2020/01/20/politics/dershowitz-impeachment-much-more-correct-cnn/tv/index.html>.

Thus desirability can be gradable even if fittingness is not.

However, there is a major problem with Maguire’s proposal: it doesn’t work for thick fittingness properties in which the act or attitude at issue is not gradable. Shall we say that x is more memorable than y when it is fitting to “more remember” x than y ? That x is more forgivable than y when it is fitting to “more forgive” someone for x than for y ? That x is more seaworthy than y when it is fitting to “more sail” x than y ? That x is more choiceworthy than y when it is fitting to “more choose” x than y ? Remembering, forgiving, sailing, and choosing are not gradable, so we cannot apply Maguire’s proposal in these cases. But memorability, forgivability, seaworthiness, and choiceworthiness are all gradable properties.²⁷

A second problem with Maguire’s suggestion is that it is not plausible that properties with an analysis of the form

x is F \equiv_{df} it is G to H x

are gradable in terms of the embedded act or attitude of H-ing, even when H-ing is gradable. Suppose we are playing a game in which we are allowed to be (pretend) angry at various people for doing various things, and we introduce a new term, ‘anger-allowed’, such that:

x is anger-allowed \equiv_{df} it is allowed for us to be angry at x .

Does it follow that *being anger-allowed* is a gradable property whose comparative form is given as follows?

x is more anger-allowed than y \equiv_{df} it is allowed for us to be more angry at x than at y .

No. Maybe the rules for how angry we are allowed to be at someone, when we are allowed to be angry at that person, have nothing to do with the rules specifying whom we are allowed to be angry at. (Perhaps I am allowed to be angry at you only when you step on a crack in the sidewalk, but how angry I am allowed to be varies with your age.) In fact, *even if* the rules for how angry we are allowed to be at someone are intimately related to the rules for whether we are allowed to be angry at that person, the proposed definition of *being more anger-allowed* is unsatisfactory. Suppose we are allowed to be angry at someone playing the game if they step on at least three cracks, and the more cracks they step on, the greater a level of anger we are allowed to have.

²⁷ Although here I use choiceworthiness as one of my examples, I am somewhat hesitant to do so. ‘Choiceworthy’ is not a word of ordinary English but rather appears (based on my research using Google Books) to have been introduced into the language in 1847 by D. P. Chase in his translation of Aristotle’s *Nicomachean Ethics*. Since then ‘choiceworthy’ has become a part of philosopher-speak, but I think we should be wary of the term, and I find that philosophers often wobble between different ways of understanding what it comes to.

If Gertrude steps on four cracks and Ichabod steps on five, our anger at Ichabod can be greater than our anger at Gertrude, but that fact about our permitted levels of anger is not what makes it the case that Ichabod is “more anger-allowed” than Gertrude, if indeed he is—and I would say that he isn’t, because *being anger-allowed* is not even a gradable property, much less a gradable property whose comparative form is determined in the manner specified.²⁸ But Maguire’s proposal predicts that it *is* a gradable property and that its comparative form *is* determined in precisely that manner.²⁹

Thus we should reject Maguire’s proposal. And given the failure of Maguire’s attempt to explain how thick fittingness properties can be gradable while the thin ones are not, the natural conclusion to draw is that it is because the thin fittingness properties are gradable that the thick ones are as well: for x to be more desirable than y is not for it to be on/off fitting to have more desire for x than for y , but rather for it to be more fitting to on/off desire x than to on/off desire y . Thus we may infer that the thin fittingness properties are gradable after all, despite the equivocal nature of the linguistic evidence for the gradability of the corresponding terms. By contrast, it is indisputable that the deontic categories *being permitted* and *being required* are non-gradable. So we have another crucial difference.³⁰

Finally, a fifth difference is that whereas permittedness and requiredness are alternatives dependent, fittingness is not alternatives dependent. Whether an act/attitude is permitted or required depends on how that act/attitude stacks up in comparison to its alternatives. But we find no such thing for fittingness: whether something is fitting (or appropriate, or merited, or . . .) depends on the relationship between that thing and

²⁸ Note that properties determined by a threshold along some dimension are not for that reason gradable: someone with a billion dollars is not “more of a millionaire” than someone with a million dollars and one cent (Moon 2017, 771–72).

²⁹ My argument here can be supported by cross-linguistic considerations. In addition to a fittingness suffix ‘-tos’ that corresponds to ‘-able’ or ‘-worthy’ in English, ancient Greek contains a requirement suffix ‘-teos’ that has no direct analogue in English. So not only is there the word ‘hairetos’ meaning ‘choiceworthy’ or ‘that which it is fitting to choose’, but there is also the word ‘haireteon’ meaning ‘that which ought to be chosen’. However, whereas ‘hairetos’ is gradable, ‘haireteon’ is not. Thus ancient Greek appears to assume that it is the gradability of the applied normative category that determines the gradability of each of these adjectives, not the gradability of the act or attitude to which the normative category is applied. (My thanks to Jacob Rosen for the point.)

³⁰ Does our argument for *fitting*’s gradability extend to *correct* as well, so that it turns out *correct* is gradable, and hence Dershowitz did not misspeak after all? I find that difficult to swallow. One solution would be to insist that although some of the thin fittingness properties are gradable, others of them are not, and it is the gradable ones that are featured in Brandt-style analyses. A second solution is to claim that ‘correct’ is what linguists call a “maximum standards absolute adjective,” on the model of ‘pure’ (Kennedy 2007): just as we cannot have two pure samples, one of which is purer than the other, perhaps we cannot have two correct claims, one of which is more correct than the other. But although that helps explain the infelicity of Dershowitz’s statement, it doesn’t explain why we are also reluctant to say that one incorrect claim is “more correct” than another incorrect claim. A third solution is to deny that correctness is a thin fittingness property. None of these solutions strikes me as entirely satisfactory, and I am not entirely sure what to say about this issue.

the situation, not on how that thing stacks up in comparison to its alternatives. Recall our guiding metaphor: whether the key fits the lock is determined by the relation between the lock and the key, not the relation between the lock, the key, and the other keys that could possibly be used instead. I do not need to consider the normatively relevant properties of the alternatives to admiring some trait of yours in order to determine that this trait of yours merits admiration and hence is admirable; all I need to consider is that trait and how it relates to an attitude of admiration. Similarly, when we assess whether a job candidate is hireable, we assess whether they are suitable for hiring independently of how hiring them would compare to hiring one of the other candidates in the pool instead. Fittingness, suitability, appropriateness, aptness: these are not alternatives-dependent notions.³¹

5. Fittingness Is Not Evaluative

So thin fittingness is not a deontic category; in particular, it is not a type of permittedness or requiredness. Could thin fittingness instead be an evaluative category; in particular, could it be a type of goodness?

Notice, off the bat, that this would be a very surprising conclusion. Although a number of authors claim that certain thick fittingness properties (such as *being admirable* and *being disgusting*) are evaluative properties,³² the claim that the thin fittingness properties (such as *being fitting* and *being warranted*) are evaluative properties is pretty much never made—and for good reason, I would say. I can think of at least five reasons to doubt that thin fittingness is a form of goodness.

First, fittingness and goodness have different guiding metaphors. Recall the guiding metaphors behind fittingness: a puzzle piece locking into place and a key fitting into a lock. These are not natural metaphors to invoke in the context of evaluative thought and talk. When consequentialists claim that certain

³¹ What about meritedness? When we claim that some award is merited, aren't we making a comparison between all the people or products or performances that are up for the award? Most likely yes, but that isn't the comparison that matters when we are assessing whether meritedness is alternatives dependent: what matters is whether we are comparing the normatively relevant properties of *giving the award to x* to the normatively relevant properties of *giving the award to y instead*, and that latter comparison is not one we are making.

³² For instance, D'Arms and Jacobson do so in the articles cited earlier. They sometimes offer a quick argument for this conclusion: thick fittingness properties are evaluative, D'Arms and Jacobson tell us, because thinking that something has one of these properties "is thinking it good or bad in a way" (2017, 251; I have harmlessly translated their argument in terms of concepts into an equivalent one in terms of properties). I worry, though, whether we really do think this unless we are using the phrases 'good in a way' and 'bad in a way' so capaciously that even requiredness counts as good in a way and forbiddenness as bad in a way. There are also a number of cases where neither applies: is deeming something to be surprising deeming it to be good or bad in a way? Finally, ways of being good or bad need not themselves be evaluative properties but rather in many cases are the non-evaluative grounds of evaluative properties: having no ink left is a way for a pen to be bad, but having no ink left is not itself an evaluative property.

states of affairs are good in themselves, I tend to imagine a cosmic thumb's up next to one of those states of affairs, or maybe a happy face hovering above it. What I don't imagine is a key fitting into a lock or a puzzle piece snapping into place, and I trust that I am not idiosyncratic in this regard. But, as before, this is a warm-up reason that is not probative on its own.

Second, if fittingness is a form of goodness, then fitting-attitude theories of value have been entirely misconceived by their proponents. This research program is almost always pitched as a way of analyzing (or reducing, or explaining) goodness in terms of other normative categories. But if fittingness is a form of goodness, then fitting-attitude theories of value end up analyzing (or reducing, or explaining) one form of goodness in terms of another form of goodness. This is another reason that is not probative on its own, since philosophers can misconceive the nature of their research programs, but it should give us pause.

Third, and more significantly, whereas goodness has a non-privative, polar opposite (namely, badness), fittingness does not have a non-privative, polar opposite. 'Unfitting' means 'not fitting' (or, perhaps, 'not fitting, and the type of thing that could be fitting'), not 'anti-fitting'. So unfittingness is merely a lack of fittingness, not the presence of anti-fittingness, whatever that might be. Similarly, for something to be undeserved is for it to be non-deserved, not for it to be anti-deserved, and for something to be unmerited is for it to be non-merited, not for it to be anti-merited.³³

Fourth, and equally significantly, there is no property that stands to *being fitting* and *being unfitting* as *being neutral* stands to *being good* and *being bad*. As we noted earlier, there is a neutral status that lies in between *being good* and *being bad*: some outcomes are neither good nor bad but in between, some knives are neither good nor bad but in between, and so on. But there is no neutral status in between *being fitting* and *being unfitting*: it is not the case that some emotional reactions are neither fitting nor unfitting but in between, that some awards are neither merited nor unmerited but in between, that some beliefs are neither justified nor unjustified but in between, and so on.³⁴

³³ Moreover, it is not just that we lack a term for the polar opposite of each of these thin fittingness properties, but rather we have no grasp on what such a thing would be, so inventing a new term—'anti-fitting', 'inverse-merited', etc.—will not help.

³⁴ Perhaps it might be objected that there are, in fact, some cases in which a fit-evaluable act or attitude is not fitting and also not unfitting. Probably the best candidate here is a desire for something that is neutral in value, such as a desire that there be parsley on the moon (Nagel 1970, 45) or that the total number of atoms in the universe be prime (Kagan 1998, 37). Could such desires be neither fitting nor unfitting? No, for at least two reasons. First, everyone must agree that an *incredibly strong* desire for a neutral state of affairs is unfitting, but it is altogether unclear how we can explain this result if moderate and weak desires for neutral things are neither fitting

Fifth, whereas goodness is gradable, it is at least arguable that fittingness is not gradable. Earlier we saw that there is some room, however slight, for arguing that the thin fittingness properties—especially *being correct*—are not gradable, even though we ultimately concluded that they are in fact gradable. However, if the thin fittingness properties are forms of goodness, and we know this, then it should be completely obvious that they are gradable. So the equivocal nature of the evidence for the gradability of the fitting, the appropriate, the proper, the correct, and the like gives us some reason to doubt that these are all forms of goodness.

Thus perhaps it is understandable why thin fittingness properties are not often claimed to be evaluative in nature. But what about the claim that at least some thick fittingness properties are evaluative? Don't advocates of fitting-attitude theories of value make that very claim? And wouldn't such a claim subvert my assertion that about the distinctiveness of the fittingness categories?

No, it wouldn't. If anything, fitting-attitude theories of value threaten the distinctiveness of the evaluative, not the fitting, for at most such analyses show that evaluative categories are a type of fittingness category, not that fittingness categories are a type of evaluative category. Moreover, I doubt they show even that. Just as we do best to interpret maximizing consequentialism as a thesis about how a certain sort of bestness is explanatorily prior to a certain sort of requiredness, where this priority relation requires the two categories to be distinct, I think we do best to interpret fitting-attitude theories of value as being concerned with an explanatory relation between two separate categories: a type of fittingness with regard to attitudes and a type of value with regard to objects. For fitting-attitude theorists, it is *because* certain attitudes are fitting that a given thing has the value it does, where that 'because' is irreflexive.³⁵ Fitting-attitude theories of value

nor unfitting. If some moderate or weak desires lie in the grey zone between fittingness and unfittingness because their object's neutrality gives them nothing to match or anti-match, why does ramping up the desire's strength eventually make this fail to be the case? (One is tempted to say that then the desire's strength no longer fits the neutrality of its object, but to say this is to go back on the claim that the desire lacked both fittingness and unfittingness before its strength was increased.) Second, a proposal on which desiring the neutral is neither fitting nor unfitting is naturally paired with a view on which blaming someone for something neutral is neither fitting nor unfitting, admiring someone for something neutral is neither fitting nor unfitting, being amused by a neither-good-nor-bad joke is neither fitting nor unfitting, and so on. But these latter verdicts are beyond the pale: if I admire you for having decided to tie your right shoe before your left one, surely my admiration is unfitting (unwarranted, unmerited, undeserved: pick your term). So a proposal that ties a desire's lack of fittingness or unfittingness to the neutrality of its object predicts that we will find many more cases in which something is neither fitting nor unfitting than we can plausibly claim to find. And more generally, any proposal that gives up on the idea that unfittingness is simply non-fittingness is left with the difficult task of explaining why it is so hard to find a single plausible (non-vague) case in which blame is neither fitting nor unfitting, admiration is neither fitting nor unfitting, etc. For these reasons, I think it is better to hold that instead of the cases under discussion being ones in which we have some neutral thing, *n*, such that desiring *n* is neither fitting nor unfitting, instead what we have is the following: desiring *n* is unfitting, but being adverse to *n* is also unfitting.

³⁵ Furthermore, even if I am wrong that we should interpret fitting-attitude theories of value in this way, so that instead they should be understood as proposing something stronger such as an identity between properties, concepts, or word meanings, recall

are compatible with the evaluative being distinct from the fitting, in precisely the way in which maximizing consequentialism is compatible with the deontic being distinct from the evaluative.

6. Going Too Far?

I have just argued for the hypothesis that there are at least three separate families of normative categories: the deontic, the evaluative, and the fitting. Some of the central differences we have found between the core categories in each family are as follows:

	<i>Required/permitted /forbidden/unrequired</i>	<i>Good/neutral/bad</i>	<i>Fitting/unfitting</i>
Duality?	Yes	No	No
Polar opposition?	No	Yes	No
Gradable?	No	Yes	Probably yes
Neutral state?	No	Yes	No
Alternatives dependent?	Yes	No	No

In a way, my argument here is overkill: any one definitive difference between each pair of families is enough to establish their distinctness. Moreover, there are distinguishing features we have uncovered that are not captured by this chart; for example, one of our primary reasons for rejecting the identification of fittingness with either permittedness or requiredness was the difficulty we encountered choosing between these two options.

So I think the overall case for my hypothesis that the fittingness categories are their own special group of normative categories is strong. However, I recognize that this hypothesis faces some significant challenges. One sort of challenge involves pointing to canonical deontic or evaluative categories that it would appear I am committed to classifying as fittingness categories, as a way of showing that my way of dividing up the normative landscape has gone too far. For instance, I can imagine an objector replying to my tripartite

that it was open to maximizing consequentialists to interpret their position in that stronger way as well, yet no one would for that reason stop distinguishing the deontic from the evaluative. So we still have the following claim, which in the end is what I am most concerned to establish: we have as much reason to distinguish the fitting from both the deontic and the evaluative as we have to distinguish the deontic from the evaluative in the first place.

division between the deontic, the evaluative, and the fitting as follows: “What about *right* and *wrong*? If *proper* and *correct* are thin fittingness properties, then surely *right* is a thin fittingness property as well. But *right* is a canonical deontic property! So something must have gone wrong somewhere.”

I have two options for how to reply to this challenge, either of which strike me as a fine way to go. The first is to deny that rightness is a thin fittingness property (perhaps by also denying that correctness is a thin fittingness property). In defense of such a move, it is striking that the opposite of ‘right’ in English is ‘wrong’, whereas every other thin fittingness term’s opposite is formed in English by adding a negative prefix (‘unfitting’, ‘inapt’, ‘unmerited’, ‘improper’, etc.). So something different is going on with rightness and wrongness, at least linguistically if not conceptually and metaphysically.

The second option, which I prefer, is to accept that *being right* is a thin fittingness property but to hold that it was a mistake for people to assume that rightness and wrongness are deontic categories. So on this proposal—which I recognize is *extremely* controversial—‘wrong’ and ‘forbidden’ are not synonyms, and ‘right’ and ‘wrong’ are not deontic terms. That might sound like crazy talk, but over time I have come to think that this suggestion is not as outlandish as it might initially seem. For one thing, it is telling that authors who assume rightness to be a deontic category sometimes get confused about whether rightness is a form of permission or of requirement.³⁶ Moreover, although authors occasionally tell us that ‘right’ is ambiguous between a reading on which it means ‘required’ and a reading on which it means ‘permitted’ (a semantic thesis for which, as far as I can tell, there is no linguistic evidence), these authors never make an analogous distinction between a reading of ‘wrong’ on which it means ‘forbidden’ and a reading on which it means ‘unrequired’, which is a puzzling gap given that wrongness is rightness’s complement. It is also telling that although we can distinguish the monadic property *being wrong* from the dyadic relation *wronging someone*, there is no analogous distinction for *being forbidden*: there is no directed notion of *doing something forbidden with regard to someone*; but if ‘wrong’ and ‘forbidden’ were synonyms, it would be puzzling why this would be so. Finally, there is strong linguistic evidence that we do not use ‘right’ as a deontic term in everyday parlance, even when we mean to be invoking a distinctively moral form of assessment. Is “Do the right thing!” just a way of saying

³⁶ My favorite example of this: in the preface to *Principia Ethica*, G. E. Moore equates “a right action” with “a duty” (1903, iv), but presumably it is more plausible to equate a right action with a permitted action (Ross 1930, 3–4).

“Do the permitted thing!” for some suitable form of permission? When I tell someone, “You were quite right to criticize him in the way you did,” is this equivalent to saying, “You were quite permitted to criticize him in the way you did”? Once we take off our moral-philosopher hats and look afresh at how we use the term ‘right’ on a daily basis, even in moral contexts, I think the idea that rightness is not a deontic category becomes quite plausible—or, at least, not as outrageously implausible as it might seem at first blush.

So that is how I would handle one sort of going-too-far challenge. Here is another challenge of that sort: “What about permissibility? Isn’t *being permissible* amenable to the same Brandt-style analysis that you endorsed for other normative ‘-ible’/‘-able’ adjectives in English, in that for ϕ -ing to be permissible is for it to be fitting to permit ϕ -ing? But that would make permissibility a thick fittingness property, whereas you yourself took it to be a canonical deontic property. So something has gone wrong.” This is another important challenge; indeed, it is probably the going-too-far worry that most concerns me. Here my preference is to deny that the property in question is in fact a fittingness property. As we use the term now, ‘permissible’ does not mean ‘fit to be permitted’, and the property it picks out is not a property that concerns the fittingness of a certain act, the act of permitting someone to do (or think, or feel) something. To say that ϕ -ing is permissible is to normatively assess *ϕ -ing itself*, not to normatively assess *the granting of a permission to ϕ* . Thus ‘-ible’ in ‘permissible’ is not functioning as a fittingness suffix as it does in ‘contemptible’ or ‘credible’. It might seem that I am making a bold linguistic hypothesis here, but in fact ‘permissible’ behaves quite differently from standard thick fittingness terms. For instance, ‘contemptible’, ‘credible’, and the other ‘-ible’/‘-able’ fittingness terms are gradable, but ‘permissible’ is not. Similarly, for most thick fittingness terms it is possible to construct a nearly synonymous term using the suffix ‘-worthy’, but not so for ‘permissible’: ‘contempt-worthy’ is close in meaning to ‘contemptible’, ‘belief-worthy’ close to ‘credible’, ‘admiration-worthy’ close to ‘admirable’, etc., but ‘permission-worthy’ is jarringly far in meaning from ‘permissible’. Although almost every normative adjective ending in ‘-ible’ or ‘-able’ is a thick fittingness term, ‘permissible’ is a rare exception.

7. Other Families?

My focus so far has been on establishing that there is a third family of normative categories beyond the most familiar two. A natural question arises, though: are there yet more families other than these? Two families

immediately spring to mind. First, there are the aretaic categories: those concerning virtue and vice. Second, there are the reason-related categories: those concerning reasons, reason, and reasoning.³⁷ I want to set aside the aretaic categories; whether they constitute a fourth family is too large a topic to settle here. However, it is worth considering the reason-related categories, because they might be thought to give rise to an objection to what I have argued so far.

This objection can be spun in one of two ways. First, it might be claimed that the reason-related categories are a third family of normative categories beyond the deontic and the evaluative which have been widely studied for several decades now, and that the fittingness categories are best identified with some subset of the reason-related categories; for instance, maybe being fitting is the same as being supported by a certain sort of reason. Thus, the objection goes, my thesis about the distinctness of fittingness as a normative category only appears surprising when we focus exclusively on the deontic and evaluative categories and ignore the reason-related categories. A second way of spinning this objection is not to claim that the reason-related categories are a separate family outside the arc of the deontic and the evaluative, but rather to claim that the reason-related categories are a subset of the deontic categories which I have neglected to mention so far. Thus, the objection continues, my arguments that fittingness is not deontic were incomplete, because those arguments were directed against the possibility that fittingness is a form of permission or requirement, not against the possibility that it is a way of being reason supported.

Both of these ways of spinning the objection rely on an assumption I deny: namely, that the reason-related categories are a unified family of normative categories. In fact the reason-related categories crosscut the various families of categories we have considered so far. Some of them are straightforward deontic categories of the sort we have already been discussing. For instance, *there being sufficient reason to ϕ* is a type of permittedness, and *there being decisive reason to ϕ* is a type of requiredness; after all, the two are duals of each other, are non-gradable, are alternatives dependent, etc.³⁸ Other reason-related categories are fittingness

³⁷ A third further family that springs to mind is what some—including my past self (Berker 2013, 382, n. 64)—call, following Michael J. Zimmerman (2002, 554), ‘the hypological categories’: those categories relating to responsibility, blameworthiness, and praiseworthiness. I now prefer to think of the hypological as a tiny and not particularly cohesive subset of the fitting.

³⁸ Thus my arguments against fittingness being a form of permittedness (or requiredness) are also arguments against analyzing fittingness in terms of what there is sufficient (or decisive) reason of the right kind to do.

categories of a kind already mentioned. For instance, *being reasonable* when it is applied to acts and attitudes (as opposed to, say, persons) is a thick fittingness property of the standard ‘-ible’/‘-able’ sort. Finally, there are favoring relations such as *being a reason to* and *having a reason to*.

Some authors classify these last two relations as deontic categories, but I do not see why they should be. Perhaps the thought here is that these favoring relations are just *pro tanto* versions of requiredness, which is certainly something suggested by the somewhat older terminology of ‘*pro tanto* duties’ (as a way of updating W. D. Ross’s unfortunate phrase ‘*prima facie* duties’ so as to avoid the suggestion that what is at issue is whether certain things are “at first glance” duties). But conceiving of reasons that way is confused, and I think it is a good thing that philosophers have largely moved away from the terminology of *pro tanto* duties and obligations. ‘*Pro tanto*’ literally means ‘as far as that goes’, so, taken literally, a *pro tanto* duty is something that considered on its own would be a duty: when no other *pro tanto* duties are present, it is one’s duty *sans phrase*, as Ross used to say. But if that is what a *pro tanto* duty is, then the category should not be gradable, for it cannot be the case that one thing is more-a-duty-*sans-phrase*-when-considered-on-its-own than another. However, the *being a reason to* relation is gradable, so ‘*pro tanto* duty’ is a misleading label for it, and we cannot conclude that this relation is a deontic category on the basis of such terminology. Even if it is true that a reason when present on its own determines one’s overall duty,³⁹ that is not what a reason is, but at most something that follows from what a reason is (together, perhaps, with what a duty is).

On an alternate and by now more common way of understanding the label ‘*pro tanto*’, we should not take that label literally but instead should understand it as a synonym for ‘contributory’. Hence the practice of referring to a consideration as “a *pro tanto* reason to ϕ ”: this phrase cannot mean “as far as that goes a reason to ϕ ” (or else we must already have a handle on the sort of thing it is when present on its own, so what work is the ‘*pro tanto*’-qualifier doing?), but rather must mean something like “a contributory reason to ϕ .” But contributes toward what? If we mean “contributes toward ϕ -ing’s requiredness,” then perhaps we can deem the relation in question to be a deontic category. However, that is another confused way of conceiving of

³⁹ And even that is in doubt: for instance, it does not appear to be true of epistemic reasons for belief (when one has a very weak epistemic reason to believe that p and no other epistemic considerations bear on the matter, one is not epistemically required to believe that p), and it might be that there are reasons that can only be present when other reasons are present (Dancy 2004, 19–20).

what reasons are. ϕ -ing's requiredness is not the sort of thing that a reason can contribute toward or count in favor of. A reason to ϕ favors ϕ -ing, not ϕ 's requiredness: it is acts and attitudes *themselves* that are favored, not facts about their normative status. Talk of "contributory reasons to ϕ " should be understood as shorthand for "considerations that contribute toward the case in favor of ϕ -ing," and it is not clear why reasons so construed should be viewed as deontic in nature. As a substantive normative matter, it might well be that what one is required or ought to do is determined by the overall balance of contributory reasons for and against the act in question and its alternatives, but this is no more reason to deem the reason relation to be a deontic category than the consequentialist's claim that what one is required or ought to do is determined by the comparative goodness of the resulting outcomes is a reason to deem goodness to be a deontic category.

So I would deny that *being a reason to* is a deontic relation. Indeed, it appears, if anything, to have all the properties we have found evaluative categories to have. *Being a reason to* lacks a dual. It has a polar opposite: namely, *being a reason against*. Both it and its polar opposite are gradable: one consideration can favor some act more than another does, and a given fact can disfavor some attitude more than another does.⁴⁰ There is a neutral relation that lies in between *being a reason to* and *being a reason against*: namely, the relation a consideration bears to an act or attitude when it counts neither in favor of nor against that act or attitude. Finally, *being a reason to* is alternatives independent: whether r is a (perhaps outweighed) reason to ϕ doesn't depend on the normatively relevant properties of the alternatives to r and to ϕ -ing.

These features of the *being a reason to* relation mean that in the end it doesn't much matter for my purposes whether that relation qualifies as deontic, because I can take my arguments against fittingness being a form of goodness and redeploy them as arguments against the suggestion that something's being fitting is the same as its having some particular type of *pro tanto* reason in its favor. If that is what fittingness comes to, we would expect it to have a polar opposite; but it doesn't. We would also expect there to be a neutral status

⁴⁰ In taking these relations to be gradable, we must be cautious. In English, the locution 'is a reason to' is *not* gradable: although we speak of reasons having a certain feature—namely, their weight or strength—that is gradable (as in: "That reason is stronger than this one"), such talk is not the same as saying that reasons *themselves* are gradable. In order to conclude that reasons are gradable, we must make the further claim that a reason is not something separate from its weight but rather consists in that weight in a given direction and nothing more, in the way in which a vector is not something separate from its magnitude but rather consists in that magnitude in a given direction and nothing more. However, this additional claim does seem plausible, especially given the widespread view that " r is a reason to ϕ " and " r favors ϕ -ing" are two ways of saying the same thing—which is naturally paired with a view on which " r_1 is a stronger reason to ϕ than r_2 is" and " r_1 favors ϕ -ing more than r_2 does" are also two ways of saying the same thing.

the lies in between fittingness and its opposite; but we find no such thing. Finally, it should be completely obvious that fittingness is gradable; but it isn't.⁴¹ In short, *being fitting* has a variety of structural features that distinguish it from *being reason supported*.⁴² So even if I am wrong when I insist that *pro tanto* reasons are not a deontic category, it still would not follow that fittingness is a deontic category because it is a certain way of being favored by a certain sort of *pro tanto* reason.

8. Conclusion

Philosophy is like all disciplines: once we have a theoretical framework that works well for us in one area, we have an unfortunate habit of trying to apply that framework everywhere else, even when it doesn't fit. I have been arguing that such a situation has occurred with the normative categories. We are all so familiar with the evaluative and the deontic categories that we assume the fittingness categories must conform to one or the other of these comfortable options. But we should stop trying to choose whether to cram the round peg of fittingness into either the square hole of requiredness or the triangular hole of goodness. Fittingness is its own thing, not a type of goodness or a variety of permittedness.

Moreover, fittingness (broadly construed) is everywhere, both in daily life and throughout philosophy. The appropriate and the merited, the commendable and the deplorable, the reasonable and the justified, the compelling and the suggestive, the liable and the excusable, the choiceworthy and the noteworthy—these are all species of fittingness. Fittingness is not some niche topic to be investigated only by the coterie of metanormative theorists seeking to explain everything normative in terms of the fitting, and 'fitting' is not a technical term that is the private reserve of neo-sentimentalists. Any moral theorist who thinks that the blameworthy is that which is worthy of blame and the preferable is that which it is fitting to prefer—and really, what are you doing not thinking these things, if you profess to understand the words 'blameworthy' and 'preferable'?—should be interested in the nature and normative significance of fittingness.

⁴¹ More precisely: it should be completely obvious that *either* fittingness itself is gradable *or* it has this feature, its strength or weight, that is gradable. (See previous footnote.) But this disjunction is also not obvious.

⁴² There is an additional structural problem with the claim that *being fitting* is the same as *being supported by a reason of the right kind* that has no parallel in my argument against identifying fittingness with goodness. $\langle x \text{ is supported by a right-kind reason} \rangle$ entails $\langle x \text{ is supported by a right-kind reason} \rangle$, but $\langle x \text{ is fitting to some extent} \rangle$ does not entail $\langle x \text{ is fitting} \rangle$: that's why a person who is admirable to some extent need not be admirable full stop, a story that is amusing to some extent need not be amusing full stop, etc. And if we try to address this problem by identifying *being fitting* not with *being supported by a reason of the right kind* but instead with *being supported by a sufficiently strong reason (or collection of reasons) of the right kind*, then *being fitting* will no longer be gradable.

Taking the fitting to be a distinct family of normative categories leaves many questions about those categories unsettled. But it also frees us up to see those categories for what they really are when we theorize about them and in terms of them, and to appreciate the ways in these categories are particularly well suited for certain domains where many other normative categories seem out of place—the realm of aesthetics and the realm of the emotions being two obvious examples. Fittingness, meritedness, aptness, and the like do not need to earn their keep by being reduced to value, or oughtness, or reasons, and perhaps it is precisely because they cannot be so reduced that they are the right categories to reach for when normatively assessing, say, aesthetic experiences or emotional reactions.

I opened by noting how the two more familiar families of normative categories are usually labeled ‘the right’ and ‘the good’, especially when moral theorists ask that perennial question, “Is the good prior to the right, or the right prior to the good?” Although most of us became acquainted with this question through the work of John Rawls, it was used as a way of structuring the theoretical options in ethics—as a way of dividing up the terrain between what we now call ‘consequentialists’ and what we now call ‘deontologists’—for most of the twentieth century. It is also a question that we can ask in normative disciplines other than ethics: for example, in epistemology we can ask, “Is the (epistemically) good prior to the (epistemically) right, or the (epistemically) right prior to the (epistemically) good?” However, if my arguments in this chapter are correct, we need to revise the way in which this traditional question is phrased. What is at issue are the priority relations between *three* families of categories, not two. Moreover, since ‘right’ might end up being a fittingness term, we are better off labeling the deontic categories as ‘the required’, not ‘the right’. So instead of asking, “Is the good prior to the right, or the right prior to the good?” what we should be asking is, “What are the priority relations between the required, the good, and fitting?”

References:

- Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, Mass.: Harvard University Press.
- Aristotle. 1847. *The Nicomachean Ethics*, translated by D. P. Chase. London: Whittaker and Co.
- Berker, Selim. 2013. “Epistemic Teleology and the Separateness of Propositions.” *Philosophical Review* 122: 337–93.
- _____. 2018. “The Unity of Grounding.” *Mind* 127: 729–77.

- _____. 2019. "The Explanatory Ambitions of Moral Principles." *Noûs* 53: 904–36.
- Brandt, Richard. 1946. "Moral Valuation." *Ethics* 56: 106–21.
- Bykvist, Krister. 2009. "No Good Fit: Why the Fitting Attitude Analysis of Value Fails." *Mind* 118: 1–30.
- Chuard, Philippe, and Nicholas Southwood. 2009. "Epistemic Norms without Voluntary Control." *Noûs* 43: 599–632.
- Cushman, Fiery, Liane Young, and Marc Hauser. 2006. "The Role of Conscious Reasoning and Intuition in Moral Judgment: Testing Three Principles of Harm." *Psychological Science* 17: 1082–89.
- Dancy, Jonathan. 2004. *Ethics without Principles*. Oxford: Oxford University Press.
- Danielsson, Sven, and Jonas Olson. 2007. "Brentano and the Buck-Passers." *Mind* 116: 511–22.
- D'Arms, Justin. 2005. "Two Arguments for Sentimentalism." *Philosophical Issues* 15: 1–21.
- D'Arms, Justin, and Daniel Jacobson. 2000a. "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research* 61: 65–90.
- _____. 2000b. "Sentiment and Value." *Ethics* 110: 722–48.
- _____. 2006a. "Anthropocentric Constraints on Human Value." *Oxford Studies in Metaethics* 1: 99–126.
- _____. 2006b. "Sensibility Theory and Projectivism." In *The Oxford Handbook of Ethical Theory*, edited by David Copp (Oxford: Oxford University Press), 186–218.
- _____. 2017. "Whither Sentimentalism? On Fear, the Fearsome, and the Dangerous." In *Ethical Sentimentalism: New Perspectives*, edited by Remy Debes and Karsten R. Stueber (Cambridge: Cambridge University Press), 250–67.
- Feldman, Richard. 2008. "Modest Deontology in Epistemology." *Synthese* 161: 339–55.
- Frick, Johann. MS. "Dilemmas, Luck, and the Two Faces of Morality."
- Gertken, Jan, and Benjamin Kiesewetter. 2017. "The Right and the Wrong Kind of Reasons." *Philosophy Compass* 12/5: 1–14.
- Gibbard, Allan. 1992. "Thick Concepts and Warrant for Feelings." *Proceedings of the Aristotelian Society, Supplementary Volumes* 66: 267–83.
- Hieronymi, Pamela. 2008. "Responsibility for Believing." *Synthese* 161: 357–73.
- Howard, Christopher. 2018. "Fittingness." *Philosophy Compass* 13/11: 1–14.
- Jacobson, Daniel. 2011. "Fitting Attitude Theories of Value." *Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2011/entries/fitting-attitude-theories/>.
- Kagan, Shelly. 1998. *Normative Ethics*. Boulder, Colo.: Westview Press.
- Kendi, Ibram X. 2019. *How to Be an Antiracist*. New York: Random House.

- Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45.
- Kjellmer, Göran. 1986. "Legible but Not Readable: On the Semantics of English Adjectives in *-ble*." *Studia Neophilologica* 58: 11–38.
- Kyle, Brent G. 2020. "The Expansion View of Thick Concepts." *Noûs* 54: 914–44.
- McHugh, Conor. 2017. "Attitudinal Control." *Synthese* 194: 2745–62.
- McNamara, Paul. 2019. "Deontic Logic." *Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2019/entries/logic-deontic/>.
- Moon, Andrew. 2017. "Beliefs Do Not Come in Degrees." *Canadian Journal of Philosophy* 47: 760–78.
- Moore, G. E. 1903. *Principia Ethica*. London: Cambridge University Press.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Princeton, New Jersey: Princeton University Press.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391–423.
- Rosen, Gideon. 2015. "Real Definition." *Analytic Philosophy* 56: 189–209.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, Mass.: Harvard University Press.
- Sennet, Adam. 2016. "Ambiguity." *Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2016/entries/ambiguity/>.
- Shoemaker, David. 2017. "Response-Dependent Responsibility: or, A Funny Thing Happened on the Way to Blame." *Philosophical Review* 126: 481–27.
- Smith, Michael. 2005. "Meta-ethics." In *Oxford Handbook of Contemporary Philosophy*, edited by Frank Jackson and Michael Smith, 3–29. Oxford: Oxford University Press.
- Suikkanen, Jussi. 2009. "Buck-Passing Accounts of Value." *Philosophy Compass* 4/5: 768–79.
- Tappolet, Christine. 2004. "Through Thick and Thin: *Good* and Its Determinates." *Dialectica* 58: 207–21.
- _____. 2013. "Evaluative vs. Deontic Concepts." In *International Encyclopedia of Ethics*, edited by Hugh LaFollette, 1791–99. Oxford: Wiley-Blackwell.
- Väyrynen, Pekka. 2021. "Thick Ethical Concepts." *Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2021/entries/thick-ethical-concepts/>.
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, Mass.: Harvard University Press.
- Zimmerman, Michael J. 2002. "Taking Luck Seriously." *Journal of Philosophy* 99: 553–76.