

IS THERE ANTI-FITTINGNESS?

SELIM BERKER
Harvard University

The permissible and the forbidden are privative opposites: each is a lack of the other. The good and the bad are, by contrast, polar opposites: badness is anti-goodness, not non-goodness. What about the fitting and the unfitting, the appropriate and the inappropriate, the apt and the inapt, the warranted and the unwarranted? Is unfittingness non-fittingness or anti-fittingness, inappropriateness non-appropriateness or anti-appropriateness? This essay argues that each of these “aptic” categories stands in a privative rather than a polar relation to its opposite. More generally, there is no coherent notion of anti-fittingness, no inversely charged flipside to aptness, to be found. In order to establish these claims, a taxonomy of different types of oppositeness is proposed, and several tests for distinguishing distinct varieties of opposites are developed. What emerges is a better appreciation of the structural characteristics of fittingness and the other aptic categories, as well as an argument for taking up the nature of oppositeness as a serious philosophical topic that is ripe for further exploration.

1. Introduction

Fittingness has been widely discussed of late. But there are many fundamental aspects of fittingness that remain underexplored. This essay focuses on one such aspect: in addition to a given response being fitting (or merited, or warranted), can a response be anti-fitting (or anti-merited, or anti-warranted)? Is unfittingness the polar opposite of fittingness, in the way that badness is the polar opposite of goodness and pain the polar opposite of pleasure, or is unfittingness a privative opposite and hence nothing more than a lack of fittingness? Settling these issues teaches us something not just about the nature of fittingness, but also, more generally, about the nature of opposition.

Contact: Selim Berker <sberker@fas.harvard.edu>

2. Forms of Fit

Fittingness has come into prominence as a result of partially overlapping work in two subdisciplines: the philosophy of emotion, where some authors—self-styled “neo-sentimentalists”—have defended the idea that affective states such as amusement and anger can be assessed for their fittingness in addition to being evaluated for their moral import or prudential upshot (D’Arms & Jacobson 2000a; 2000b; 2023); and metanormative studies, where some authors—so-called “fittingness firsters”—have posited that fittingness is the fundamental normative category in terms of which all other normative categories can be explained (Chappell 2012; Howard 2019; McHugh & Way 2016, 2022). However, fittingness should be of interest to all scholars working in normative areas of research, not just neo-sentimentalists or fittingness firsters and their respective critics, for the fittingness categories suffuse our thinking about normative matters. The fitting and the proper, the apt and the appropriate, the merited and the deserved, the warranted and the justified, that which is called for and that which is to be done: these are all forms of fittingness, broadly conceived. So too are the preferable and the deplorable, the credible and the contemptible, the useful and the shameful, the surprising and the boring, the persuasive and the repulsive, the trustworthy and the blameworthy. Fittingness has been right in front of us all along, and if anything, what “calls for explanation” (to use yet another fittingness locution) is not why fittingness has become such an active area of research lately, but rather why it took so long for fittingness to come into focus as its own distinctive subject.

My own hypothesis is that fittingness has not been given its due because of the prominence of the question “Is the good prior to the right, or the right prior to the good?” as a way of thinking about the structure of normative theories. Implicit in this question is the assumption that there are two overarching families of normative categories—evaluative categories such as *good*, *bad*, *better*, *best* (“the good”) and deontic categories such as *ought*, *permitted*, *forbidden* (“the right”)—whose comparative explanatory priority can then be explored. This assumption then tempts us into thinking that all normative categories must be either evaluative or deontic. So we are led to the view that some fittingness categories—for example, the warranted and the justified—must be deontic categories, whereas others—for example, the admirable and the boring—must be evaluative categories, because when we are given only two options, those seem to be the most likely respective candidates. In this way we assimilate the fitting into the deontic and the evaluative, and thereby overlook its uniqueness, as well as its ubiquity.

It will be useful to have an adjective meaning ‘pertaining to fittingness (and the like),’ in the way that ‘deontic’ means ‘pertaining to duty (and the like)’ and ‘evaluative’ means ‘pertaining to value (and the like).’ I propose the following:

'apptic.'¹ In my view (Berker 2022), the apptic categories are a third major family of normative categories which are as different from the evaluative and the deontic categories as those two families of normative categories are from each other.² But my project in this essay does not rest on that view of mine. Even those who think of the apptic categories as a subset of the deontic and/or evaluative categories should be interested in their structural features and fundamental nature. Recall just how prevalent these categories are both within the philosophy classroom and outside of it. Accountability, blameworthiness, credibility, desirability, eligibility—these are all apptic notions. So too are the attractive, the bothersome, the commendable, the disgraceful, and the exhilarating.

We can distinguish three layers of apptic categories. At the most basic level, there is a relation of fit that obtains between an object, *O*, and a response, *R*, just in case *R* fits *O*, as when our admiration fits a philosopher's accomplishments, or my disgust fits the rancid food, or your blame fits my transgression. This relation is not symmetric: when *R* fits *O*, it need not follow that *O* fits *R*, in the relevant sense.³ The converse of fitness we pick out with a variety of locutions: when *R* fits *O*, we say that *O* is worthy of *R*, that *O* calls for *R*, that *O* merits *R*, or that *O* warrants *R*. So not only does your blame fit my transgression, but my transgression warrants (or merits, or calls for) blame on your part—these are two sides of the same credit card. Here I use 'response' to refer to *the sorts of entities that can bear relations of fit* and 'object' to refer to *the sorts of entities toward which relations of fit can be born* without taking a stand on the exact range of entities that fall under each; for instance, perhaps failing to have some attitude or emotion counts as a

1. Many thanks to Britta Clark for suggesting the term. Two other candidates I considered were 'axiotic,' from the Greek 'axios' meaning 'worthy' or 'merited' (but often mistranslated as 'having value,' due to a conflation of fittingness and goodness), and 'hypological,' from the Greek 'hypologos' meaning 'liable or held responsible.' The first of these I rejected because the term 'axiology' is so firmly entrenched as meaning 'the study of value' (when in fact it should mean 'the study of worthiness or meritedness'). And the second I rejected because Michael J. Zimmerman originally introduced 'hypological' to mean "having to do with moral responsibility" (2006: 585; see also 2002: 554), but I want a term that encompasses a broader range of fittingness notions, and also because I want to avoid having to refer to the study of fittingness via the abomination 'hypology' (whereas 'apptology' rolls off the tongue).

2. Here I use 'normative' as a catch-all that encompasses the deontic, the evaluative, the apptic, and anything else of those sorts. This usage has become standard in much of contemporary philosophy, but there are parts of the English-speaking philosophical world where some authors use 'normative' to mean what I mean by 'deontic.' To these philosophers, I ask, "So what is your term for something that is either 'normative' (in your sense) or evaluative or apptic? Wouldn't it be useful to have a term which covers that entire disjunction?" Since 'normative' literally means 'pertaining to standards or norms,' and since there can be standards of goodness and of fittingness in addition to standards of oughtness, 'normative' is well placed to serve as this general catch-all term—and not, it might be added, well placed to serve as a synonymy for 'deontic.'

3. For that reason, matching metaphors—such as those employed in (Hurka 2022); (Lee 2022); and (McHugh and Way 2022: 77)—are ill-suited to pick out relations of fit, since matching is a symmetric relation.

“response” in the relevant sense, and perhaps the relevant “object” that some response fits is a situation or set of facts rather than a spatiotemporal object that can be tripped over or knocked into.

So the first layer of aptic categories involves dyadic relations of fit between responses and objects (and those relations’ converses). The second layer consists in monadic properties born by responses when they fit the object to which they are a response. When your blame fits my transgression (a two-place relation), then my blame is fitting (a one-place property).⁴ When a philosopher’s body of work merits our admiration (a two-place relation), then our admiration is merited (a one-place property). Let us call these monadic properties *the thin aptic properties*. Beside the fitting and the merited, other examples include the warranted, the called for, the suitable, the justified, the deserved, the proper, the appropriate, and—of course—the apt. Many of these monadic properties are picked out in English by adjectives that are derived from verbs denoting a corresponding dyadic relation at our first layer: so ‘suitable’ is derived from ‘suits,’ ‘justified’ from ‘justifies,’ ‘deserved’ from ‘deserves,’ and so on. But some of our terms in English for thin aptic properties were inherited from other languages without the verbs from which they were derived being passed along as well, as in the case of ‘apt,’ ‘appropriate,’ and ‘proper.’⁵

Finally, our third layer of aptic categories is by far the most extensive. It consists in a range of normative properties that are picked out in English by adjectives ending in distinctively aptic uses of the suffixes ‘-able’/‘-ible,’ ‘-worthy,’ ‘-ful,’ ‘-ing,’ ‘-ive,’ and ‘-some,’ among others.⁶ As was pointed out by Richard Brandt (1946), these properties apply to an object when it is fitting to have a response of some designated sort to that object (or, equivalently, when the object merits, calls for, or is worthy of that sort of response). So we have:

- For *O* to be admirable is for it to be fitting to admire *O*.
- For *O* to be credible is for it to be fitting to believe *O*.
- For *O* to be trustworthy is for it to be fitting to trust *O*.
- For *O* to be shameful is for it to be fitting to feel ashamed of *O*.
- For *O* to be surprising is for it to be fitting to be surprised by *O*.
- For *O* to be repulsive is for it to be fitting to be repulsed by *O*.
- For *O* to be bothersome is for it to be fitting to be bothered by *O*.

4. Thus we need to distinguish fitness from fittingness: the former is a relation, the latter a property. (Or, if fittingness is really an *n*-place relation whose additional relata are often left implicit in ordinary conversation, then fittingness is an *n*-place relation and fitness an (*n*+1)-place relation.)

5. For example, ‘apt’ comes from a Latin adjective derived from the Latin verb ‘apere,’ meaning ‘to fasten.’

6. Other suffixes with aptic uses: ‘-y’ (as in ‘funny’), ‘-ous’ (as in ‘outrageous’), and perhaps ‘-ant’/‘-ent’ (as in ‘important,’ ‘salient’).

Aptic uses of ‘-able’/‘-ible’ are particularly prevalent in English: once you start noticing them, it is hard to stop. Although a few aptic suffixes (e.g. ‘-ful’) are no longer productive in contemporary English, most of them are, and new words employing these suffixes are readily created.⁷ Aptic terms ending in ‘-able,’ ‘-ible,’ ‘-worthy,’ and ‘-ful’ usually involve a response to an object that is expressed in the active voice (when *O* is admirable, what is fitting is that I admire *O*), whereas those ending in ‘-ing,’ ‘-ive,’ and ‘-some’ almost always involve a response that is expressed in the passive voice (when *O* is surprising, what is fitting is that I am surprised by *O*, or, equivalently, that *O* surprises me). Let us call the properties designated by these terms *the thick aptic properties*.

There are delicate questions about how exactly each of these layers relates to each other. The Brandt-style analysis of thick aptic properties in terms of thin ones endorsed above is a template whose details can be filled out in various ways that I shall not settle here.⁸ I shall also not pursue the question of how exactly the monadic properties at our second layer are built up out of their corresponding dyadic relations at the first layer. These issues are worthy of further exploration, but pursuing them is not my primary task. We need not settle every question about the aptic categories in order to make progress on some. Compare: we need not settle how *pro tanto* reasons combine so as to determine what there is overall reason to do in order to make progress in theorizing about the *pro tanto* reason relation.

Instead, I want to focus on the neglected question of how negation works among the aptic categories. All of the positive properties at our second layer have corresponding negative properties: in addition to the fitting, the apt, the appropriate, the merited, the warranted, etc., there is the unfitting, the inapt, the inappropriate, the unmerited, the unwarranted, etc. But what exactly is the relation between the fitting and the unfitting? Are they privative opposites, as we find with the permissible and the impermissible, so that one is merely the absence of the other? Or are they polar opposites, as we find with goodness and badness? Badness is not simply a lack of goodness but rather is the inverse of goodness—anti-goodness, as it were, not non-goodness. Is unfittingness non-fittingness, or is it anti-fittingness? Similarly, is unmeritedness non-meritedness or anti-meritedness? Is unwarrantedness non-warrantedness or anti-warrantedness? We can also ask parallel questions about the thick aptic properties at

7. Consider ‘dissertable,’ meaning ‘fit to be written about in a Ph.D. dissertation’ (as used in Diamond 2019: 100), and ‘grammable,’ meaning ‘fit to be photographed for Instagram’ (as used in Srinivasan 2020).

8. For example, when *O* is admirable, is it fitting for everyone to admire *O*, or only certain suitably specified agents? Is it fitting for the agents in question to admire *O* on any basis whatsoever, and to any degree, and for any length of time, or are there constraints on the type of admiration and its temporal dimensions? For discussion, see (Achs & Na’aman 2023).

our third layer. Is the undesirable the non-desirable or the anti-desirable? Is the untrustworthy the non-trustworthy or the anti-trustworthy?

Answering these questions is crucial to our understanding of the aptic categories and has direct bearing on how we theorize about them. If fittingness has a polar opposite, then it is not enough to provide an account of the fittingness conditions for a given emotion; one also owes us an account of the anti-fittingness conditions for that emotion, since each set of conditions will not simply be the complement of the other.⁹ Moreover, these theories of the conditions under which an emotion is fitting or anti-fitting will be constrained in various ways: each must appeal to the sort of thing that has a polar opposite, given the plausible principle that if (a) being F and being G are polar opposites, and (b) being F^* fully makes it the case that a thing is F , then it must be the case that (c) being F^* has a polar opposite, being G^* , such that being G^* fully makes it the case that a thing is G . If fittingness has a polar opposite, we cannot take a given type of emotion to be made fitting by a condition without a polar opposite. Similar constraints will apply to theories of the conditions under which a belief is justified, if there is such a thing as anti-justifiedness, and to theories of the conditions under which a distribution of benefits and burdens is deserved, if there is such a thing as anti-deservedness.

But before we can directly assess whether aptic categories such as fittingness have polar opposites, we should clarify what polar opposition is and how it can be distinguished from other sorts of oppositeness relations.

3. Varieties of Opposition

Oppositeness (or opposition, in the logical sense of that term) is a surprisingly neglected topic in philosophy. Linguists have extensively investigated the phenomenon as it arises in natural language and proposed a variety of typologies for lexical opposition relations, but it is difficult to find a single systematic study of opposition relations among contemporary philosophers. So it is largely to linguists that we must turn when seeking guidance for how to categorize different types of oppositeness.

Our topic here is not exactly the same as that of linguists. We are interested in opposition relations borne by categories (properties, relations, entities, etc.), not by the words we use to pick out those categories. When we say that ignorance is the opposite of knowledge, or aversion the opposite of desire, or giving

9. Compare the literature on ill-being (i.e. negative well-being), where it is insisted that the conditions that make for ill-being are not simply the logical complement of the conditions that make for well-being (Kagan 2014; Rice 2019; Sumner 2020; Bradford 2021, 2022). I take it the underlying assumption is that ill-being is well-being's polar, not privative, opposite.

the opposite of receiving, we are talking about the things themselves, not about words in English such as ‘ignorance,’ ‘knowledge,’ and so on. But in carefully thinking about how certain pieces of language are opposites of each other, linguists have uncovered many helpful insights.

Indeed, it can even be argued that linguists who study oppositeness of meaning—or “antonymy,” as they call it—are already theorizing about a metaphysical relation of opposition, despite their claims to be primarily studying a semantic relation.¹⁰ The argument goes as follows. Two words (or other lexical items) are said to be antonyms when they have opposite meanings. So if words W_1 and W_2 are antonyms, this is in virtue of (i) W_1 ’s having meaning M_1 , (ii) W_2 ’s having meaning M_2 , and (iii) M_1 ’s being the opposite of M_2 . But M_1 and M_2 are not themselves words; instead they are properties, or functions from possible worlds to extensions, or the mental states expressed by atomic sentences employing those words, or something else of that sort—of course the details will vary depending on one’s theory of meaning and the grammatical type of W_1 and W_2 . The key point is that the oppositeness relation between M_1 and M_2 is not a relation between lexical items but rather a more fundamental relation of opposition between non-lexical items that bear the meaning relation to lexical items. So when thinking about whether W_1 and W_2 are antonyms and what the various features of this relation of antonymy are, we can only do so by thinking about whether M_1 and M_2 are opposites and what the features of this more fundamental oppositeness relation are. You can’t know that the words ‘happy’ and ‘sad’ are antonyms without knowing that the property *being happy* is the opposite of the property *being sad*, or that a function from possible worlds to the happy people in those worlds is the opposite of a function from possible worlds to the sad people in those worlds, or what have you. In short, linguists have already been thinking about a metaphysical relation of opposition among categories when theorizing about antonymy and other semantic forms of opposition among lexical items.¹¹

10. As in the title of M. Lynne Murphy’s *Semantic Relations and the Lexicon*, which devotes an entire chapter to antonymy and claims it to be “the archetypical lexical semantic relation” (Murphy 2003: 169).

11. Here is another way of making the point. Antonymy—oppositeness of meaning—is supposed to contrast with synonymy—sameness of meaning. But when words W_1 and W_2 are synonyms, it is not the words themselves that stand in a sameness (or identity) relation, but rather something else, namely their meanings. So, too, when W_1 and W_2 are antonyms, it is not the words themselves that stand in an oppositeness relation, but rather something else, namely their meanings.

(Incidentally, thinking about this parallel reveals another point. Although it is often said that antonymy is the opposite of synonymy, this common dictum is strictly speaking false. The opposite of sameness is distinctness, not oppositeness. So one would expect the opposite of synonymy to be non-synonymy, not antonymy. ‘Antonym’ and ‘synonym’ are not antonyms.)

In what follows, I sketch my own preferred taxonomy of oppositeness relations. It builds on some influential taxonomies that one finds in the linguistics literature but also deviates from them in several important respects. After noting some of these differences, I turn to the issue that shall be of most interest to us: How does one determine the type of oppositeness relation at issue in cases where one is confident that the two categories are opposites, but unsure of the type? I propose several tests for distinguishing opposition relations from one another that we shall go on to use when turning to our main topic, namely the varieties of oppositeness that one finds among the aptic categories.

I find it useful to distinguish three fundamentally different types of oppositeness. First, there is *privative oppositeness*. When categories *A* and *B* are privative opposites, then at least one of them is the absence or lack of the other. Sometimes this absence-of relation goes in both directions: forbiddenness is a lack of permissibility, and permissibility is a lack of forbiddenness. Other times the absence-of relation appears to go in only one direction: sobriety is a lack of drunkenness, but we do not usually think of drunkenness as a lack of sobriety. Since lacking a property is the same as not having that property, properties that are privative opposites can be characterized equally well by way of logical negation rather than by using ‘lack’- or ‘absence’-talk: to be forbidden is to be not permitted; to be permitted is to be not forbidden; to be sober is to be not drunk. Other pairs of categories plausibly classified as privative opposites include *alive/dead*, *married/single*, *healthy/sick*, *clean/dirty*, *partial/impartial*, *afraid/unafraid*, and *with/without*.

Second, there is *polar oppositeness*. When categories *A* and *B* are polar opposites, each is the inversely charged flipside of the other, or its metaphysical mirror image, as it were. Here I resort to metaphors, but I do so unapologetically: with our most fundamental metaphysical categories we often have no recourse other than to lean on metaphors when characterizing them. Positive and negative charge are two paradigm polar opposites (and what I mean to be invoking with the label ‘polar,’ rather than the North and South Pole). Two other paradigm polar opposites are goodness and badness. In both cases, neither item in each pair is simply the lack of the other: some chemicals are neither positively nor negatively charged, some outcomes neither good nor bad. Instead, each side of a pair of polar opposites is its own positive thing (‘positive’ in the metaphysical sense of not being a lack, not ‘positive’ in an electrochemical or normative sense). Other pairs of categories plausibly classified as polar opposites include *pleasure/pain*, *happy/sad*, *beautiful/ugly*, *approve/disapprove*, *friend/enemy*, *inside/outside*, *clockwise/counterclockwise*, and *for/against*.

Third, there is *conversive oppositeness*. Direct conversive opposites, such as *borrow from/lend to*, consist in two relations that are converses of each other, so that their subject and object slots have been swapped, whereas indirect conver-

sive opposites, such as *borrower/lender*, consist in two categories that are suitably related to two relations that are converses of each other. Other pairs of categories plausibly classified as converse opposites include *buy/sell*, *give/receive*, *win/lose*, *parent/child*, *teacher/student*, and *above/below*.

I bring up this third type of oppositeness only to set it aside. Converse opposites are puzzling. On the one hand, including a category of converse oppositeness in our taxonomy seems forced upon us in order to account for paradigm opposites such as *give/receive* and *parent/child*. On the other hand, not every relation and its converse are plausibly opposites of each other (blaming and being blamed are not opposites, nor are seeing and being seen), and it is altogether unclear why we deem some converses to be opposites and others not to be. Moreover, whereas there is no overlap between the categories of privative and polar opposition,¹² it does appear to be possible for a pair of converse opposites to also be privative or polar opposites. For instance, a case can be made that *win/lose* are privative opposites (in virtue of winning being not losing and vice versa) in addition to being converse opposites (in virtue of their connection to the relation *beats* and its converse), and *above/below* are plausibly both polar and converse opposites (if *up/down* are polar opposites, surely *above/below* are as well). So perhaps all putative converse opposites are really privative or polar opposites that just happen to involve a relation and its converse, or a suitable connection to such a relation. Either way, fittingness and the other aptic categories at our second and third layers clearly do not have converse opposites, so for our purposes we need not settle whether converse opposition is genuinely distinct from privative and polar opposition.

Two comments before we proceed. First, notice that privative opposition relies on the notion of a relevant domain in which the privative opposites apply. For a thing to be sober is for that thing to lack drunkenness—but only when ‘a thing’ ranges over creatures that can be intoxicated. That is why the number 7 is not sober, despite its failing to be drunk. Similarly, it is not true that any entity whatsoever that fails to be married is single, or that fails to be alive is dead, or that fails to be forbidden is permitted. In each case, the relevant lack or absence is restricted to a certain domain.¹³ This restriction to relevant domains is not unique to privative opposition: we shall see that there is a way in which polar opposites

12. The one possible exception: maybe *presence/absence* are both privative and polar opposites, if we deem the absence of absence to be presence (and/or the absence of presence to be absence), and also deem absence to be the metaphysical flipside of presence. I am inclined to reject the latter claim: absence is not its own (ontically) positive thing since it is, well, an absence of a thing.

13. Exactly which domain? Often that will be a difficult question to answer. But our uncertainty about the exact boundaries of the relevant domain is no impediment to our classifying two categories as privative opposites. We might have trouble deciding whether the Pope is within the domain of people to which the categories *being single* and *being married* apply, but that does not make us question whether those two categories really are privative opposites.

pick out a relevant domain as well. Second, our discussion of converse opposites introduced the notion of direct and indirect opposition. When two categories, *A* and *B*, are indirect opposites, this is because they can be defined in terms of other categories that are themselves (direct or indirect) opposites. In the simplest case, *A* is defined in terms of *A**, *B* is defined in terms of *B**, and *A** and *B** are opposites. But not any old way of defining *A* and *B* in terms of *A** and *B**, respectively, suffices to make *A* and *B* indirect opposites. Perhaps *bringing about a benefit* is the indirect polar opposite of *bringing about a burden*, in virtue of benefits and burdens being polar opposites and the bringing-about relation being an oppositeness-transmitting relation. But *being to the north of a benefit* is not the opposite of *being to the north of a burden*. It is an interesting topic for future research exactly which forms of definition transmit oppositeness and why.

My tripartite division among fundamental types of opposition is roughly similar to the most common taxonomies of opposition that one finds in the linguistics literature: most include a grouping roughly corresponding to my notion of privative opposition, a second grouping roughly corresponding to my notion of polar opposition, and a third leftover grouping (or set of groupings) for the odds and ends not covered by the first two.¹⁴ One important difference is that whereas linguists focus almost exclusively on pairs of adjectives or verbs that constitute opposites, and hence on pairs of properties and relations once we translate their lexical taxonomies into a corresponding metaphysical taxonomy, I have imposed no restrictions on the types of categories that can stand in oppositeness relations. Although it is natural to focus on pairs of opposites that can be picked out using adjectives and verbs (or, more generally, using adjectival and verbal phrases), we should not forget that many paradigm opposites fall into categories that correspond in natural language with adverbs (e.g. *quickly/slowly*), nouns (e.g. *benefit/burden*), prepositions (e.g. *for/against*), determiners (e.g. *some/none*), and many more.

A second important difference is that I do not define privative and polar opposition in terms of complementarity and contrariety. Many linguists draw a taxonomic distinction between the sort of opposition we find when two predicates—call them '*F*' and '*G*'—are *complements* (or *contradictories*) and the sort we find when they are *contraries*. For '*F*' and '*G*' to be complements is for both of the following to be the case:

- a. No OVERLAP: Necessarily, nothing in the relevant domain is both *F* and *G* (and so, necessarily, everything that is *F* is not *G*, and everything that is *G* is not *F*).

14. For representative examples, see (Lyons 1968; 1977), (Palmer 1976), (Cruse 1976; 1980; 1986; 2002), (Lehrer & Lehrer 1982), (Lehrer 1985; 2002), (Murphy 2003), and (Gao & Zheng 2014).

- b. No GAP: Necessarily, nothing in the relevant domain is neither *F* nor *G* (and so, necessarily, everything that is not *F* is *G*, and everything that is not *G* is *F*),

For '*F*' and '*G*' to be contraries is for instead the following to be the case:

- a. No OVERLAP: Necessarily, nothing in the relevant domain is both *F* and *G* (and so, necessarily, everything that is *F* is not *G*, and everything that is *G* is not *F*).
- b*. MAYBE GAP: Possibly, something in the relevant domain is neither *F* nor *G*.

Now it is true that, when the categories whose oppositeness is at issue are both properties, then those two categories will be complements of each other whenever they are privative opposites, and they will be contraries of each other whenever they are polar opposites. But when we are dealing with opposites that are categories other than properties (or relations), defining what complementarity and contrariety come to is no trivial matter. (In what sense are burdens the contrary of benefits, and for-ness the contrary of against-ness?) Moreover, taking contrariety to be definitive of polar opposition is particularly problematic, because many contraries are not plausibly opposites: for example, *being neutral in value* and *being good* are not opposites despite being contraries, nor are *being prime* and *being divisible by 12*. Contrariety is at most a necessary condition for polar opposition, not a sufficient condition. So to use the label 'contrary opposition' for this type of opposition is to mislead, since what we have is a type of opposition that when present leads to contrariety, not a type of opposition defined by the presence of contrariety. 'Polar' is my attempt at capturing what I take to be definitive of this variety of opposition. Similarly, complementarity—the dividing of a domain into two exclusive and exhaustive sets—is in my view at most a consequence of the sort of opposition evinced by *dead/alive* and *married/single*, not what it fundamentally comes to, which I take to be its privative aspect—its connection to an absence or a lack.¹⁵

A third important difference between my taxonomy of oppositeness relations and those one encounters in the linguistics literature is that I do not tie any of my types of oppositeness to gradability. One of the most influential taxonomies of opposition in linguistics over the past half century is that of John Lyons (1968: §10.4; 1977: §§9.1–9.3). Lyons draws a tripartite distinction between (1) the sort of oppositeness exemplified by *married/single*, which—under the assumption that

15. Note that I follow linguists in using 'privative' to mean 'pertaining to an absence or a lack,' not to mean 'pertaining to a privation': no assumption is being made that the absent thing is normally present, or that its absence constitutes a defect.

these come to the same thing—he calls both ‘ungradable oppositeness’ and ‘complementarity’; (2) the sort of oppositeness exemplified by *good/bad*, which he calls both ‘gradable oppositeness’ and ‘antonymy’; and (3) the sort of oppositeness exemplified by *borrow/lend*, which he calls ‘converseness.’ Lyons considers type (2) to be oppositeness “*par excellence*” (1968: 463), and for that reason he restricts the term ‘antonymy’ so that it covers only type (2) oppositeness and not types (1) and (3) as well—a bizarre terminological convention that many linguists working on lexical oppositeness still follow to this day.¹⁶ To make matters worse, Lyons’ very taxonomy rests on an assumption that his terms ‘ungradable opposites’ and ‘complements’ are co-extensive, but it was quickly pointed out that this is not so, for there exist so-called “gradable complementaries” such as *clean/dirty* that are gradable (one thing can be cleaner or dirtier than another) despite being complements (if something in the relevant domain is not clean, then it is dirty, and vice versa) (Palmer 1976: 81; Cruse 1980; 1986: 203–204). This revelation left it unclear how to extend Lyons’ taxonomy to take into account gradable complementaries: do we take complementarity versus contrariety to be what distinguishes (1) from (2), so that gradable complementaries are an additional variety of type (1) oppositeness, or do we take gradability versus non-gradability to be what distinguishes (1) from (2), so that gradable complementaries are an additional form of antonymy to be added to group (2)? More generally, the pride of place that Lyons bestowed on gradability in his taxonomy has led to a trend whereby linguists both take gradable opposites to the paradigm of opposition and talk as if there is this thing, “gradable antonymy,” that can be contrasted with other forms of antonymy and oppositeness (Jones 2002; Lehrer 2002; Gao & Zheng 2014).

I think this focus on gradability is a mistake. Being gradable is not a feature of opposition relations themselves; at most it is a feature of some of the things that stand in opposition relations. Moreover, there is no distinctive form of opposition, I claim, that arises only for gradable categories or only for non-gradable categories: on my view, the gradable versus non-gradable distinction crosscuts the most fundamental differences that we find among the oppositeness relations. Although the most obvious privative opposites are non-gradable (*alive/dead*, *married/single*, *prime/composite*, etc.), the aforementioned “gradable complementaries” are all gradable privative opposites; these include *healthy/sick*, *clean/dirty*, *partial/impartial*, *afraid/unafraid*, and—as I shall argue below—a large number of thick aptic properties and their negations. Similarly, although the most obvious polar opposites are gradable (*good/bad*, *pleasure/pain*, *beautiful/ugly*, etc.), many polar opposites are non-gradable and hence do not admit of comparative forms, such as *inside/outside* (one of my possessions cannot be “more inside” my suitcase than another), *clockwise/counterclockwise* (I cannot turn my body in a “more

16. For example, all of the works cited in fn 14 adhere to this convention.

clockwise" direction than you turn turns), and *positive/negative* when applied to integers (-7 is not "more negative" than -2). Finally, although most converse opposites are non-gradable (*win/lose*, *parent/child*, etc.), there is nothing in the notion itself that rules out the possibility of converse opposites being gradable.

A fourth important difference is that I have sought to avoid any "mince-pie oppositions" in my taxonomy, so I have left off several varieties of opposition in my taxonomy that linguists include but which strike me as not *fundamental* types of opposition. The reference here is to G. E. M. Anscombe's quip that if practical reasoning were just theoretical reasoning whose conclusion concerns practical matters, then practical reasoning would be no more fundamental a type of reasoning than "mince-pie reasoning" – reasoning whose conclusion concerns mince pies – would be (1957: 58). Just as we would not include mince-pie reasoning in any reasonable taxonomy of varieties of reasoning, so too should we avoid including in our taxonomy of opposition relations any variety of opposition that just happens to be about a given subject matter without being different from other forms of opposition *qua* type of opposition, rather than *qua* type of thing opposed. One such example is so-called "reversive opposition" of the sort exemplified by *tie/untie*, *damage/repair*, and *ascend/descend*, which are all said to involve "either changes in opposite directions between two terminal states . . . or the causation of such changes" (Cruse 2002: 507). No doubt there are reasons why, from a linguist's perspective, the words picking out these sorts of opposites have interestingly different features from the words picking out other sorts of opposites. But for our purposes, there is no need to distinguish a separate category of reversive opposition, for all reversive opposites just are polar opposites when the things opposed happen to involve a process of some sort, in the way that mince-pie reasoning happens to be about mince pies. For similar reasons, we need not include a separate grouping for "directional opposition" of the sort exemplified by *up/down*, *north/south*, and *arrive/depart*, as Lyons (1977: 281–287) and Alan Cruse (1986: 223–231) do in their influential taxonomies. Directional opposition is simply polar opposition when the things opposed happen to involve contrary motions or directions.

I have just distinguished three distinct types of opposition: privative opposition of the sort exemplified by *married/single*, polar opposition of the sort exemplified by *happy/sad*, and converse opposition of the sort exemplified by *parent/child*. Now we must ask: how do we tell these types of opposition apart in cases in which we are confident that two categories are opposites but unsure of the type of opposition at issue? In particular, since we are setting aside converse opposition as not relevant to our central topic, how do we tell the difference between privative and polar opposition?

A tempting thought is that negative prefixes, when present, can be our guide. Isn't a category picked out by an English word starting with a nega-

tive prefix such as ‘in-,’ ‘un-,’ or ‘dis-’ always the privative opposite of a category picked out by an English word shorn of that prefix? As it turns out, no. Although ‘in-’ (as well as its allomorphs ‘il-,’ ‘im-,’ and ‘ir-’) often functions as an absence-maker that generates a word for the base category’s privative opposite, as we find with ‘incomplete,’ other times ‘in-’ functions as a polarity-flipper that generates a word for the base category’s polar opposite, as we find with ‘inconsiderate.’ (To be inconsiderate is not merely to fail to be considerate, but rather to actively act in ways that are the reverse of considerate behavior). Similarly, although ‘un-’ most often functions as an absence-maker, as in ‘unmarried’ or ‘unafraid,’ it sometimes acts as a polarity-flipper, as in ‘unhappy’ or ‘unprofessional.’ And although ‘dis-’ is almost always a polarity-flipper, as in ‘disapprove’ or ‘disfavor,’ there are occasions when it is more plausibly an absence-maker, as in ‘disorganized’ or ‘disinterest.’ These English prefixes do not on their own distinguish privative from polar opposition.¹⁷

But there are other negative prefixes that do unambiguously pick out either privative or polar opposition, and it will be useful to employ these in what follows. In English, ‘non-’ always acts as an absence-maker and never as a polarity-flipper, and ‘anti-’ sometimes acts as a polarity-flipper and never as an absence-maker.¹⁸ So these prefixes are a handy way of distinguishing privative from polar opposition. To be unmarried is to be non-married, not anti-married. And to be unhappy is to be anti-happy, not non-happy. Thus in many cases, especially when we are considering categories picked out by adjectives, nouns, and verbs in natural language, it will be useful to rephrase our question of how we tell whether two categories, *A* and *B*, count as privative or polar opposites as the question of how to determine whether *B*-ness is anti-*A*-ness or non-*A*-ness.¹⁹

We can deploy a number of tests to answer this question of whether two opposites are of the privative or polar variety—that is, of the anti- or non- sort. First, there is the most straightforward test, which I call *the Direct-Intuition Test*: we tend, in many cases, to have rather firm intuitions directly on the issue of whether a given oppositeness relation is privative or polar in nature.²⁰ So when

17. The same is true of many negative affixes (prefixes, infixes, suffixes, etc.) in many languages other than English (Zimmer 1964).

18. The polarity-flipping use of ‘anti-’ (as in ‘anticlimax’ or ‘antimatter’) must be distinguished from an adversarial use (as in ‘antiaircraft’ or ‘antibacterial’). It is the former I mean to be invoking throughout.

19. Actually, my use of ‘non-’ here is imperfect. Although ‘non-’ in ordinary English sometimes picks out an absence or lack that is the base category’s privative opposite, as in ‘non-fungible,’ other times ‘non-’ picks out an absence or lack that is not necessarily restricted to some relevant domain and is not necessarily the base category’s privative opposite, as in ‘non-voluntary.’ Throughout I intend to be using ‘non-’ in this first manner.

20. I formulate this test in terms of intuitions, but skeptics about intuitions in philosophy and elsewhere can replace my appeal to a faculty of intuition with an appeal to whatever faculty we use to make judgments of oppositeness.

confronted with a pair of opposites, we can ask ourselves, “Is it plausible that one or both of these categories is by the definition the lack or absence of the other?” And we can ask ourselves, “It is plausible that each of these opposites is its own positive thing that is the metaphysical flipside of the other, in the way that badness is the metaphysical flipside of goodness and pain the metaphysical flipside of pleasure?”²¹ If we can confidently answer either of these questions in the affirmative, then we can safely sort the opposition relation as privative or polar, respectively.

But what about cases in which our intuitions about the type of opposition at issue yield no verdict, or do yield a verdict but only a rather hesitant one? For example, consider a philosophically interesting pair of opposites: *coherent/incoherent*. I am quite confident that incoherence is the opposite of coherence, but when I directly consider the matter I find myself unsure as to whether incoherence is simply a lack of coherence, or whether instead being incoherent is a negatively charged version of the positive property of being coherent. In cases such as these we need to resort to other means in order to determine the type of opposition at issue.

A second test, which is the primary one employed in the linguistics literature, I call *the Neither-Nor Test*: if properties *being F* and *being G* are polar opposites, then it is possible for there to be an item in the relevant domain that is neither *F* nor *G*, whereas if they are privative opposites this is not possible. So when we have already determined that *being F* and *being G* are opposites and narrowed down the type of opposition to either polar or privative opposition, the possibility or impossibility of finding an entity of the relevant sort that is neither *F* nor *G* settles the matter. Some people are neither happy nor unhappy (polar opposition), but no person is neither married nor unmarried (privative opposition). Some chemical compounds are neither acidic nor basic (polar opposition), but no chemical compound is neither pure nor impure (privative opposition). A similar test can be employed for categories other than properties. There are some people I neither like nor dislike (polar opposition), but there are no commands I neither obey nor disobey (privative opposition). Some people are neither for nor against a given proposal (polar opposition), but no person is neither with nor without an opinion on the matter (privative opposition).

This is a powerful and versatile test, but it comes with two important caveats. First, the restriction to a relevant domain is crucial. We do not want to say that *drunk* and *sober* fail to be privative opposites because the number 7 is neither drunk nor sober. When applying the Neither-Nor Test, we need to consider whether *the sort of thing that can be assessed as either drunk or sober* (for example, a person) can be neither drunk or sober. When the categories whose oppositeness

21. Of course here ‘positive thing’ must be interpreted ontically, not normatively.

is at issue are both gradable, one handy way of making sure we are restricting our assessment to items in the relevant domain is by way of the locution ‘in between’: we can ask, “Can something be neither *F* nor *G* but in between?” Although the number 7 is neither drunk nor sober, it is not neither-drunk-nor-sober-but-in-between. So our inability to conceive of anything that is neither drunk nor sober but in between allows us to conclude that drunkenness and sobriety are privative opposites. By contrast, we are able to find chemical compounds that are neither acidic nor basic but in between, attitudes that are neither approval nor disapproval but in between, and eventualities that are neither a benefit nor a burden but in between, so these opposites are all polar, not privative.²²

Second, we need to ignore vagueness when applying the test. It is a vague matter exactly where the borderline between a dish’s being dirty and its being clean lies, and on some (but not all) theories of vagueness this implies that there are some dishes that are neither dirty nor clean. But we still want to say that *dirty* and *clean* are privative opposites. So on these theories of vagueness, we need to apply the Neither-Nor Test to precisified versions of the categories at issue. And when we do, we get the desired result: for any given precisification of ‘dirty’ and ‘clean,’ it follows that no dish can be neither dirty nor clean. Thus, strictly speaking, the way to apply the Neither-Nor Test to properties *being F* and *being G* is to ask, “Vagueness aside, can there be something in the relevant domain that is neither *F* nor *G*?” In practice we have no trouble applying the test despite the need for this qualification given some (but, again, not all) theories of vagueness. ‘Alive’ and ‘happy’ are both vague terms, but—vagueness aside—no organism is neither alive nor dead, whereas—even when we set vagueness to one side—some people are neither happy nor sad. In what follows I take these qualifications as given.

There is a third test that we can apply only in the special case in which both of the categories whose type of oppositeness is at issue are gradable; I call it *the Independent-Gradability Test*. Polar opposites are each their own thing, metaphysically speaking, whereas privative opposites are intimately tied to each other via an absence-of relation. So when two opposites are gradable, we would expect that the grading of polar opposites happens independently of each other, whereas the grading of privative opposites is always correlated. Thus in the case of privative opposition among two properties, *being F* and *being G*, the following pair of entailment relations will hold:

22. When the opposites being assessed are not gradable, the ‘in between’-locution is infelicitous and thus causes interference during applications of the Neither-Nor Test. The number 0 is neither positive nor negative, but it is infelicitous to say, “0 is neither positive nor negative but in between.” The locution is also less helpful in cases in which, although the two opposites are gradable, it is not clear they are graded along a single scale or by way of one unified dimension of evaluation. “He’s neither a good nor a bad artist but in between” sounds awkward to my ear.

- e₁. $\langle X \text{ is more (or less) } F \text{ than } Y \rangle$ entails $\langle X \text{ is less (or more) } G \text{ than } Y \rangle$.
 e₂. $\langle X \text{ is more (or less) } G \text{ than } Y \rangle$ entails $\langle X \text{ is less (or more) } F \text{ than } Y \rangle$.²³

If one person is more (or less) drunk than a second person, then the first person is less (or more) sober than the second, and vice versa. If one sample of a chemical compound is more (or less) pure than a second sample, then the first sample is less (or more) impure than the second, and vice versa. By contrast, in the case of polar opposition among properties, one or both of (e₁) and (e₂) will be false. If one person is happier than another, it doesn't necessarily follow that the first person is less sad than the second person. And if one painting is more beautiful than another, it doesn't necessarily follow that the first painting is less ugly than the second.²⁴

Notice that all we require is that, for gradable polar opposites, one or both of (e₁) and (e₂) is false, not that both are false. Often both sets of entailments fail to hold, as in the case of *beautiful/ugly*: not only does $\langle X \text{ is more beautiful than } Y \rangle$ fail to entail $\langle X \text{ is less ugly than } Y \rangle$, but $\langle X \text{ is uglier than } Y \rangle$ fails to entail $\langle X \text{ is less beautiful than } Y \rangle$. But sometimes our entailments among comparative polar opposites only fail in one direction. Arguably, $\langle X \text{ is more unhappy than } Y \rangle$ does indeed entail $\langle X \text{ is less happy than } Y \rangle$.²⁵ In English, the word 'happy' does double duty in a way that the word 'unhappy' does not: 'happy' both picks out a monadic property and designates a gradable quantity that ranges over the entire scale of happiness and unhappiness, whereas 'unhappy' only does the former, not the latter.²⁶ However, there is still an important sense in which comparative happiness is independent of comparative unhappiness. Although more comparative happiness entails less comparative unhappiness, the reverse does not hold.²⁷

23. Here I use parentheses as a compact way of denoting two entailment claims at once. So really (e₁) is the conjunctive claim that $\langle X \text{ is more } F \text{ than } Y \rangle$ entails $\langle X \text{ is less } G \text{ than } Y \rangle$, and $\langle X \text{ is less } F \text{ than } Y \rangle$ entails $\langle X \text{ is more } G \text{ than } Y \rangle$. Moreover, because $\langle X \text{ is more } F \text{ (or } G \text{) than } Y \rangle$ both entails and is entailed by $\langle Y \text{ is less } F \text{ (or } G \text{) than } X \rangle$, (e₁) is logically equivalent to

e₁*. $\langle X \text{ is more (or less) } F \text{ than } Y \rangle$ entails $\langle Y \text{ is more (or less) } G \text{ than } X \rangle$,

and (e₂) is logically equivalent to

e₂*. $\langle X \text{ is more (or less) } G \text{ than } Y \rangle$ entails $\langle Y \text{ is more (or less) } F \text{ than } X \rangle$.

Sometimes it is easier to consider (e₁*) and (e₂*) rather than (e₁) and (e₂) when applying the Independent-Gradability Test.

24. We can apply the Independent-Gradability Test as just described only when *both* of two opposites are gradable. So the fact that $\langle X \text{ is more talented than } Y \rangle$ fails to entail $\langle X \text{ is less untalented than } Y \rangle$ does not give us reason to doubt that *talented/untalented* are privative opposites, for although *talentedness* is gradable, *untalentedness* is not.

25. It is less clear, though, that $\langle X \text{ is sadder than } Y \rangle$ entails $\langle X \text{ is less happy than } Y \rangle$, which suggests that *being sad* and *being unhappy* are not the same property and that the true polar opposite of happiness is unhappiness, not sadness.

26. As linguists put it, 'unhappy' is marked, whereas 'happy' is unmarked. For an overview, see (Lehrer 1985).

27. Some pairs of opposites observe a *pro tanto* vs. overall distinction, and when they do there is a fourth test—which I call *the Aggregation Test*—that we can use to determine whether the type of

With these additional tests in hand, we can now return to our example in which the Direct-Intuition Test produced inconclusive results. Are coherence and incoherence privative or polar opposites? My own intuitions about the matter yield a muddled verdict. But with our other tests we can triangulate in on a more definite result—one that I find surprising. The Independent-Gradability Test tells in favor of *coherent/incoherent* being private opposites, although not conclusively so. It is clear that if one narrative is more (or less) incoherent than a second narrative, then the first narrative is less (or more) coherent than the second one. I am also inclined to accept the reverse direction—so if one narrative is less (or more) coherent than a second narrative, then the first narrative is more (or less) incoherent than the second one—but I grant that this verdict is less certain. However, when it comes to the Neither-Nor Test, there is no ambiguity: here we undoubtedly are working with privative opposites. It is just not possible to find an item assessable in terms of its coherence or incoherence that is neither coherent nor incoherent but in between. If the movie script we are working on together starts out with a coherent storyline that progressively becomes less and less coherent with each subsequent revision, we do not—vagueness to one side—reach a point at which the storyline is no longer coherent but not yet incoherent. As we use these terms in everyday parlance when describing storylines, narratives, plans, strategies, and the like,

opposition at issue is either privative or polar. (I owe the central idea behind this test to Malcolm Morano.) With polar opposites, both sides of the distinction are relevant to assessments at the overall level. When determining whether, all things considered, an outcome is good or bad or neutral, we need to take into account both the respects in which the outcome is good and the respects in which it is bad. Even if overall goodness is not determined through a simplistic procedure in which one “adds up” the respects in which the outcome is good and “subtracts” the respects in which it is bad, still there is a sense in which the respects in which the outcome is good push the total verdict toward one of overall goodness and away from one of overall badness, whereas the respects in which the outcome is bad do the reverse. But with privative opposites, only one side of the distinction matters to the overall assessment. When determining whether, all things considered, a person is healthy or sick, we do not “add up” the respects in which the person is sick and “subtract” the respects in which the person is healthy in order to come to a final estimation of whether the person is healthy or sick overall, nor do we engage in a more complicated aggregation procedure that takes into account considerations on both sides of the ledger. It is not as if a person’s being extremely healthy in some respects can compensate for, or offset, the respects in which that person is sick during an assessment of their overall health. Rather, all that matters are the respects in which the person is sick. Similarly, a sample of some chemical compound is pure or impure overall in virtue of the respects, if any, in which it is impure, not in virtue of the respects in which it is impure “minus” the respects in which it is pure, and a third-party observer is partial or impartial overall in virtue of the respects, if any, in which they are partial, not in virtue of the respects in which they are partial “minus” the respects in which they are impartial. In what follows, I do not lean heavily on this fourth test, because it is controversial whether the aptic categories obey a *pro tanto* vs. overall distinction, and I do not wish to wade into that debate here. I do believe, though, that a proper application of the Aggregation Test offers further support for the conclusions I draw in this essay.

'coherent' and 'incoherent' are privative opposites—although that is compatible with philosophers using these words as technical terms that pick out polar opposites.

4. Opposition among the Thick Aptic Categories

We can now turn to our main topic: the types of opposition found among the aptic categories. Let us start by considering thick aptic properties of the sort picked out in English via suffixes such as '-able'/'-ible,' '-ing,' '-ive,' '-some,' and '-worthy.' Do these have polar or privative opposites?

A few clearly have polar opposites. Desirability and undesirability, for instance, pass all of our tests for polar opposition with flying colors. Intuitively, the undesirable is the anti-desirable, not the non-desirable. When we consider the range of items that can be assessed for desirability or undesirability, we find some that are neither positively desirable nor negatively undesirable but rather occupy a neutral zone in between. Finally, if two outcomes are both desirable and the first is more desirable than the second, it doesn't follow that the second is more undesirable than the first, so comparative undesirability floats free from comparative desirability. Thus according to the Direct-Intuition, Neither-Nor, and Independent-Gradability Tests, the desirable and the undesirable are polar, not privative, opposites.

There are several other thick aptic categories that also clearly have polar opposites; for example, the likable and the unlikable. But overall, thick aptic categories with uncontroversial polar opposites are hard to find. Instead, what one more often encounters are (i) thick aptic properties that clearly have privative opposites, or (ii) thick aptic properties with an opposite whose status is nonobvious but on closer examination reveals itself to be privative in nature.

Consider *lovable/unlovable*, *interesting/uninteresting*, and *persuasive/unpersuasive*. These are all obvious cases of privative opposition, as our three tests confirm. Intuitively, the unlovable is the non-lovable, not the anti-lovable; the jilted lover who worries that they are unlovable is worrying that they are not fit to be loved, not worrying that they might have an inverse form of lovability (being fit to be hated?). No scientific finding is neither interesting nor uninteresting but in between; if we take a somewhat interesting finding and steadily make it less and less interesting, we do not eventually reach a point at which the finding is no longer interesting but is not yet uninteresting. If one objection in a philosophy paper is more persuasive than a second objection, then the second is more unpersuasive than the first, and vice versa; we do not have two graded categories here, comparative persuasiveness and comparative unpersuasiveness, that can somehow pull apart from each other. For similar reasons, the inexcusable is

the privative opposite of the excusable, the unconvincing the privative opposite of the convincing, the inoffensive the privative opposite of the offensive, and the blameless the privative opposite of the blameworthy.²⁸

But what about—to pick a representative difficult case—the unbearable? Here matters are less clear. Although it is perhaps more intuitive to take *bearable/unbearable* to be privative rather than polar opposites, there is some intuitive pull toward taking each side of the contrast to be its own (ontically) positive thing. So here the Direct-Intuition Test does not yield as clear a verdict in favor of privative oppositeness as it does when applied to, say, the unconvincing or the excusable. But on closer examination, I think it becomes clear that we are dealing with privative opposites in this case. The Neither-Nor Test yields particularly strong verdicts. Some jobs are bearable, others are unbearable, but if we take a job that is bearable and steadily make it less bearable, do we eventually reach a point at which the job is no longer bearable but is not yet unbearable? Surely not. The Independent-Gradability Test also yields striking results. Can we really distinguish the comparative unacceptability of two jobs from their comparative acceptability, so that the first job might be more unacceptable than the second without the first being less acceptable than the second, or vice versa? When we apply our full range of tests, it becomes hard to resist concluding that *bearable/unbearable* are indeed privative opposites.

So we have found that most thick aptic terms in English with a negative counterpart pick out privative rather than polar opposites.²⁹ Moreover, the overall pattern we have found among these terms and their corresponding properties suggests a natural hypothesis. It is striking that *desirable* and *likeable* have polar opposites, and so do the underlying responses *desiring* and *liking*, whereas, say, *excusable* and *persuasive* do not have polar opposites, and neither do the underlying responses *excusing* and *persuading*. I anti-desire something when I am averse to it, and I anti-like something when I dislike it, but I cannot anti-excuse you for some mistake or anti-persuade you of some point. So perhaps a thick aptic property has a polar rather than privative opposite when and only when the response whose fittingness is at issue itself has a polar opposite. Then we will say that the polar opposite of being *R*-worthy is being (anti-*R*)-worthy:

28. Note that ‘-less’ sometimes acts as a negative aptic suffix which functions like ‘-worthy’ and ‘un-’ applied in succession, as in ‘blameless’ (= ‘that which it is not fitting to blame or blame for’), although it more often acts as an absence-making suffix, as in ‘shameless’ (= ‘that which is without shame’). I can be blameless despite being blamed for what I did, but I cannot be shameless while being ashamed of my conduct.

29. What about thick aptic terms for which English currently lacks a prefixed (or suffixed) term denoting its referent’s opposite, such as ‘contemptible’ or ‘fearsome’? I am inclined to hold that these terms still have opposites—in these two cases, privative opposites—and it is largely an accident that English never developed terms for their opposites such as ‘incontemptible’ or ‘unfearsome’ (perhaps due to the inhibiting factors posited by Kjellmer 2005: 159–161).

For O to be R -worthy is for it to be fitting to have response R to O .
 For O to be (anti- R)-worthy is for it to be fitting to have response anti- R to O .

I mention this hypothesis because it is so natural, but fully defending it would require more space than I have here. In particular, we would need to argue that, because *lovable* does not have a polar opposite, *loving/hating* are not in fact polar opposites—a conclusion with which I am comfortable, but which I recognize others will resist.

5. Opposition among the Thin Aptic Categories

Let us now turn to a more delicate issue: whether thin aptic categories such as *fitting*, *apt*, *appropriate*, *warranted*, *merited*, and *deserved* have privative or polar opposites.

Our findings with regard to the thick aptic categories give us the resources to offer a powerful argument against the existence of a polar opposite for fittingness. If there were such a thing as anti-fittingness, then we would expect *every* thick aptic category to have a polar opposite. If the convincing is that which it is fitting to be convinced by, then the anti-convincing would be that which it is anti-fitting to be convinced by. More generally, being R -worthy would always have a polar opposite, namely being R -(anti-worthy), definable like so:

For O to be R -(anti-worthy) is for it to be anti-fitting to have response R to O .

But this is not what we find, either at the level of language or of metaphysics. English lacks a word for the polar opposite of most thick aptic properties, which would be a puzzling gap if those opposites were so plentiful. And when we considered the properties themselves, we came to the conclusion that many (if not most) thick aptic properties have a privative opposite.³⁰

One way of resisting this argument would be to insist that thick aptic properties such as the convincing and the excusable have both a privative and a polar

30. Why doesn't the existence of those few thick aptic properties with a polar but not a privative opposite cut in the opposite direction? If the unfitting is the non-fitting, shouldn't we expect every thick aptic property to have a privative opposite? No, because although an anti-version of a category is always that category's polar opposite, a lack of some category is not always that category's privative opposite: only some lacks or absences constitute opposites. (The non-good is not an additional opposite of the good, beyond the bad.) But exactly which absences constitute opposites and why is a puzzling issue that will have to be addressed in future work.

opposite, and for some reason natural language has chosen to use ‘unconvincing’ and ‘inexcusable’ to latch onto the former, not the latter. But it is just not clear that it is possible for a category to have both a privative and a polar opposite: the existence of one type of opposite seems to exclude the existence of the other type.³¹ One seeming counterexample, building on an example often mentioned in the linguistics literature,³² involves the word ‘un-American’: don’t ‘American’ and ‘un-American’ pick out polar opposites, whereas ‘American’ and ‘non-American’ pick out privative opposites? Yes, they do, but this is no threat to our claim that one and the same category cannot have both a privative and a polar opposite, because the sense of ‘American’ that contrasts with ‘un-American’ is different from the sense of ‘American’ that contrasts with ‘non-American’: an un-American news source cannot be American in this first sense, but it might well be American in the second sense.³³

A second way of resisting our inference from the types of opposition found among the thick aptic properties to the types of opposition found among the thin ones would be to insist that at most we are entitled to draw such a conclusion for those thin aptic properties that feature in Brandt-style analyses of the thick aptic properties, such as fittingness, meritedness, or aptness. But since it is not clear that, say, deservedness ever features in a Brandt-style analysis, we have left open the possibility that there exists anti-deservedness. Similarly, if there exist distinctive flavors of the aptic categories that are not plausibly insertable into Brandt-style analyses, such as moral appropriateness or epistemic warrantedness, then these too might have polar opposites, for all our initial argument has shown.

In the end I am not moved by this objection: I think it would be very peculiar if the sort of constitutive, flavorless warrantedness that plausibly features in many Brandt-style analyses lacks a polar opposite, but epistemically flavored warrantedness does have a polar opposite, and I think that deservedness does plausibly feature in a number of Brandt-style analyses—a fireable offense is an offense for which one deserves to be fired, a praiseworthy action is one that deserves praise—although it is a tricky matter to settle when and why fittingness and deservedness sometimes travel together and sometimes part ways. But there is another reason to be dissatisfied with our line of argument so far: it is very much an indirect way of trying to establish that the thin aptic properties

31. Note that this exclusion only appears to happen for privative and polar opposition. Some categories plausibly have both a converse opposite and an opposite of some other type: for instance, one opposite of *giving* is *receiving* (its converse opposite), and another is *taking* (its polar opposite). This is yet another reason why converse opposition is puzzlingly different from the other varieties.

32. See, for example, (Zimmer 1964: 33), (Quirk et al. 1985: 1540), (Horn 2020: 19), and (Joshi 2020: 82).

33. Similarly, the sense of ‘professional’ that contrasts with ‘unprofessional’ (meaning ‘anti-professional’) is different from the sense of ‘professional’ that contrasts with ‘nonprofessional.’

lack polar opposites, so it would be nice to supplement that indirect form of argumentation with a more direct consideration of how plausible or implausible it is to take the thin fittingness properties to have polar opposites. Thus let us directly consider the matter. All of the thin aptic properties have a named opposite in English picked out by a term featuring a negative prefix: 'unfitting,' 'inapt,' 'unmerited,' etc. Accordingly, I shall break up our discussion into two stages. First, I want to consider whether it is plausible that these terms pick out polar opposites of their positive counterparts. Is, for example, unfittingness anti-fittingness? Second, if it turns out that unfittingness is not anti-fittingness, I want to consider whether nevertheless there might be such a thing as anti-fittingness, although it is not picked out by the term 'unfitting,' and similarly for the other thin aptic properties. Is anti-fittingness existent but unnamed?

5.1. Is Unfittingness Anti-Fittingness?

When we consider whether unfittingness is the privative or polar opposite of fittingness, inaptness the privative or polar opposite of aptness, unmeritedness the privative or polar opposite of meritedness, and so on for all of the thin aptic properties, two of our three tests do not fully settle the matter.

For several thin aptic properties, the Direct-Intuition Test yields very strong results. It is, I would say, blindingly obvious that the unmerited is the non-merited, not the anti-merited, and the undeserved the non-deserved, not the anti-deserved. Undeserved compliments are not inverse-deserved (what could that be?), and when we consider whether an award is unmerited, that's equivalent to considering whether it is not merited. But for other thin aptic properties, the Direct-Intuition Test yields less definitive results. I myself find it more intuitive to hold that unfittingness is non-fittingness than to hold that it is anti-fittingness, but my intuitions have an uncertain edge to them here, and that is all the more so for the unjustified and the uncalled for. Moreover, a number of people report to me that they find it more intuitive to take inappropriateness to be anti-appropriateness rather than non-appropriateness. I am not quite sure what it means for inappropriate workplace behavior to be anti-appropriate workplace behavior, but I grant that our initial intuitions about the inappropriate and the unjustified are less clear than our initial intuitions about the unmerited and the undeserved. When I think about how close the unmerited is to the unwarranted, and how close the unwarranted is to the unjustified (and, equally importantly, when I think about everyday uses of the term 'unjustified,' not uses of that term in contemporary epistemology, where it has largely become a technical term), my intuitions to the effect that the unjustified is the non-justified firm up, but this sort of path-dependence to our intuitive verdicts is always worrisome.

Matters are even less clear with the Independent-Gradability Test, in part because it is controversial whether fittingness and the other thin aptic properties are gradable.³⁴ If one attitude is more unfitting than a second attitude, does it follow that the first is less fitting than the second? Yes. But if one attitude is more fitting than a second attitude, does it follow that the first is less unfitting than the second? I can see resistance to this claim, especially if it is possible for there to be two fitting attitudes, one of which is more fitting than the other, and in general I find it difficult to settle these issues just by thinking about the nature of fittingness or unfittingness themselves.

So it is to the Neither-Nor Test that we must turn when assessing whether *fitting/unfitting*, *apt/inapt*, *appropriate/inappropriate*, and the like are privative or polar opposites. And here our results are particularly clear and striking. Workplace behavior is never neither appropriate nor inappropriate but in between. Blame never falls in a grey zone in which it is both not fitting and not unfitting. Anger that fails to be apt is always inapt, and vice versa. For the vast majority of acts and attitudes that we assess by way of thin aptic properties and their named opposites, we do not allow for a gap in which neither the property nor its named opposite applies.

But does this pattern only hold in *most* cases, or does it apply to *all* fit-assessable acts and attitudes? I claim that the pattern does indeed hold with universal generality. But philosophers have proposed to me a number of alleged counterexamples to this claim. Three of them are as follows. First, consider the fittingness of desiring something of neutral value.³⁵ Some claim that such a desire is neither fitting nor unfitting. Second, consider an intention by Buridan's ass to eat the left bale of hay. Some claim that such an intention is neither fitting nor unfitting. Third, consider suspension of judgment, understood as a *sui generis* attitude that is the doxastic alternative to belief and disbelief. Some claim that such an attitude can be neither fitting nor unfitting.

All three of these would-be counterexamples are unconvincing. Start with the third one. The idea that suspension of judgment can be neither fitting nor unfitting is usually motivated by the thought that a belief is fitting iff it is true and unfitting iff it is false, but there is nothing that stands to suspension of judgment as truth or falsity stands to belief, so there are no fittingness or unfittingness conditions for suspension of judgment. A related motivation appeals to the idea that a fitting act or attitude "gets things right" (McHugh & Way 2022), but suspension of judgment is never a matter of getting things right; rather, it is a matter of being committed to not taking a stand on some question. Note, though,

34. For an argument that fittingness is non-gradable, see (Maguire 2018). For a reply, see (Berker 2022: §4).

35. Or, if you prefer, the fittingness of desiring *as an end in itself* something of neutral final value.

that both of these motivations are ways of arguing that suspension of judgment is outside of the realm of being assessable in terms of fittingness or unfittingness, not that it is inside that realm but the verdict reached is always one of being neither fitting nor unfitting. So even if suspension of judgment is never fitting or unfitting, that is no challenge to the idea that fittingness and unfittingness are privative opposites, just as the fact that natural numbers are never drunk or sober is no challenge to the idea that drunkenness and sobriety are privative opposites. What we want is an example of a *fit-assessable* act or attitude that can be neither fitting nor unfitting, and for that we have to look elsewhere.³⁶

Our second example potentially faces the same problem: it is not obvious that intentions as such are fit-assessable. (There is no natural-language term meaning ‘intention-worthy,’ nor is it easy to create one that we feel comfortable adopting.) But let us grant that intentions can, in fact, be fitting or unfitting. Then this second example might seem more promising. How can we say that intending to eat the left rather than the right bale is fitting, or vice versa? And isn’t it too harsh to say that both intentions are unfitting? But on closer inspection this example is also unconvincing. Eating the left bale and eating the right bale are not the only options available to the ass; she could also walk away (and thereby go unfed). Presumably intending to walk away is unfitting, given the circumstances. So why not say that intending to eat the left bale is fitting and intending to eat the right bale is also fitting, and indeed that they both are equally fitting? Perhaps it is unfitting for the ass to *prefer* eating the left one over eating the right one, or vice versa; but preference is not intention. While there is nothing to recommend eating the left bale over eating the right bale, that is no bar to its being the case that intending to eat the left bale and intending to eat the right bale are both fitting. Each is a fitting response to the situation, although they are not fitting in conjunction.³⁷

36. Moreover, it is just not clear that we need to accept either assumption used to motivate the idea that suspension of judgment is neither fitting or unfitting. If one’s belief that *p* is fitting not when it is true that *p*, but rather when one has sufficient evidence for $\langle p \rangle$ (or, in a variant proposal, when one is in a position to know that *p*), then we can say that suspension of judgment about whether *p* is fitting when one lacks sufficient evidence both for $\langle p \rangle$ and for $\langle \text{not-}p \rangle$ (or, in the variant, when one is neither in a position to know that *p* nor in a position to know that not-*p*). And given that it is fitting to be puzzled by puzzling comments and confused by confusing arguments, either it is false that fitting acts or attitudes always “get things right,” or this suggestion is true but no obstacle to there being a sense in which fitting suspension of judgment gets something right, namely the same sense in which fitting puzzlement or confusion gets something right. (For discussion of the potential fittingness conditions for suspension of judgment, see McHugh & Way 2022: §§5.2, 6.5; Lee 2022; Rosa 2023; Vollmer forthcoming.)

37. My diagnosis of this case becomes even more clearly the right one, I would say, when we switch from focusing on the fittingness of *intention* to focusing on the fittingness of *choice*. It is just as fitting for the ass to choose to eat the left bale as it is for the ass to choose to eat the right bale; both options are equally choiceworthy. But it would be very peculiar to say that, although both options are equally choiceworthy, neither option is choiceworthy *simpliciter*, and hence neither is fit to be chosen.

This leaves the first example: desires for things without either value or disvalue, such as an astronaut's desire to leave some parsley on the moon (Nagel 1970: 45) or an amateur mathematician's desire that the total number of atoms in the universe be prime (Kagan 1998: 37). Whims of this sort seem both common and uncriticizable.³⁸ To say that such desires have the negative status of being unfitting would appear to be unduly severe, and to say that they have the positive status of being fitting would appear to be unsustainable, given that these examples are supposed to be ones in which, by stipulation, the objects of desire are not desirable—that is, not fit to be desired.³⁹ The objects of desire in these examples are also, by stipulation, not undesirable. So they are not fit to be averse to. But there is another attitude beyond desire and aversion that fits them perfectly well—indifference.⁴⁰ How can indifference toward the possibility that the number of atoms in the universe is prime be fitting, but desire for that possibility to obtain is neither fitting nor unfitting? Desire fails to fit these various valueless objects, and an alternative to desire—indifference—does fit them, so isn't it natural to say that desiring these objects is unfitting? In saying so, we need not say that a small desire for these objects is *very* unfitting: it might only be unfitting to a small degree (assuming, for the moment, that we can measure such things in degrees). And, in fact, another advantage of this view is that then we can say that the degree of unfittingness for a neither-desirable-nor-undesirable object is proportional to the intensity of one's desire: the stronger one's desire, the more unfitting the attitude. I take it that advocates of these counterexamples do not want to hold that an incredibly strong desire for something that is neither desirable nor undesirable is neither fitting nor unfitting: surely it would be unfitting for our astronaut's overwhelming desire in life to be a matter of selenic greenery stranding. But if extremely strong desires for the neither desirable nor undesirable are unfitting, why do such desires become neither fitting nor unfitting as their intensity decreases? Not because then they better fit their object: by stipulation they don't fit their object. Because they are less of a worse fit with their object? But that would be to concede that they *are* unfitting, even when their intensity is small. I conclude that it is best to view these examples as ones in which desire for the objects in question is indeed unfitting, while keeping in mind that a small desire for those objects need not be very unfitting, and also keeping in mind that aversion to those objects is equally well unfitting.

38. Here I assume, for the sake of argument, that a whim is a special type of desire, rather than a distinct mental state that may well be distinguishable from desire precisely by its failure to be fit-assessable.

39. However, it should be noted that this stipulation is not always easily achieved. For example, it might be very slightly desirable for the astronaut to garnish the moon with some parsley, because doing so is amusing.

40. I owe this observation to Christopher Howard.

I have replied to these three types of counterexamples to my claim that no fit-assessable act or attitude is neither fitting nor unfitting because they provide a general template for how to respond to other would-be counterexamples to that claim: some putative counterexamples are cases in which the item being assessed is not fit-assessable (as we find with suspension of judgment, given certain background assumptions); other putative counterexamples are cases in which the item being assessed is in fact fitting (as we find with the ass's intentions); and yet other putative counterexamples are cases in which the item being assessed is in fact unfitting (as we find with desires for things of neutral value). Moreover, it is worth stressing again how rare even the putative counterexamples of this sort are. This rarity puts the advocates of these counterexamples in an awkward situation. If we can decrease the desirability of some object until we reach a point at which desire for that object is neither fitting nor unfitting, why is it not possible to decrease the admirability of a philosopher's accomplishments until admiring those accomplishments is neither fitting nor unfitting, and why is it not possible to decrease the persuasiveness of an argument until being persuaded by that argument is neither fitting nor unfitting? My answer: because across the board it is impossible for a fit-assessable item to be neither fitting nor unfitting, for the fitting and the unfitting are privative opposites.

5.2. *Is Anti-Fittingness Unnamed?*

Now that we have argued that unfittingness is non-fittingness, not anti-fittingness, and similarly for the other negative thin aptic categories denoted by prefixed terms in natural language, there is one last issue for us to consider: might there still be such a thing as anti-fittingness even if it is not picked out by the term 'unfitting'? And might there still be such things as anti-meritedness and anti-warrantedness even if they are not picked out by the terms 'unmerited' and 'unwarranted'? In other words, are polar opposites of the thin aptic categories existent but unnamed?

We have already seen two reasons to doubt this possibility. First, we have our indirect argument that if anti-fittingness and the like exists, then we should expect every thick aptic properties to have a polar opposite, but that is not what we find. Second, we have noted that it does not appear to be possible for an item with a privative opposite to also have a polar opposite. To these considerations, we can add a third: it just baffles the mind what anti-fittingness, anti-meritedness, or anti-warrantedness might be, if they are not unfittingness, unmeritedness, or unwarrantedness. We cannot simply create polar opposites willy-nilly. I have as difficult a time making sense of how the anti-fitting could exist beside the fitting and the unfitting as I have conceiving of how the "anti-married" could

exist beside the married and the unmarried, or how the “anti-complete” could exist beside the complete and the incomplete.

But maybe at least in the case of the anti-merited, we can make sense of this status after all. Consider the following example, which I owe to Arash Abizadeh. Suppose we are deciding which soldier in our battalion to give a medal for bravery. Soldier S_1 fought courageously in the war, at grave personal risk. Soldier S_2 acted neither courageously nor cowardly, having been assigned to a role that kept them far from the frontlines. Soldier S_3 fled from the enemy forces in a blatant display of cowardice, resulting in many casualties for our side. If these are the only three candidates, then surely giving the medal of courage to S_1 is merited, whereas giving the medal to either of the other two soldiers is not merited. However, it might be thought, there is a difference here: whereas giving the medal to S_2 is merely unmerited, giving the medal to S_3 is positively anti-merited, in addition to being unmerited. So is anti-meritedness a coherent notion after all?

I do not think so, though I grant the force of the example. Here the Independent-Gradability Test is helpful. What I want to say is that giving the medal to S_3 is less merited (and hence more unmerited) than giving it to S_2 , while denying that in addition the act has this special status, being anti-merited. The advocate of anti-meritedness presumably will make analogous comparative claims: giving the medal to S_3 is both less merited and more anti-merited than giving it to S_2 . And surely giving the medal to S_1 is even less anti-merited than giving it to S_2 . So giving the medal to S_1 is both more merited and less anti-merited than giving it to S_2 . Thus in making room for this supposed thing, anti-meritedness, we have committed ourselves to the following:

- e_1^m . $\langle X \text{ is more (or less) merited than } Y \rangle$ entails $\langle X \text{ is less (or more) anti-merited than } Y \rangle$.
- e_2^m . $\langle X \text{ is more (or less) anti-merited than } Y \rangle$ entails $\langle X \text{ is less (or more) merited than } Y \rangle$.

Hence greater meritedness is never independent of lesser anti-meritedness in either direction. But there is no other pair of gradable polar opposites for which this is so: greater beauty does not always come with lesser ugliness; more pleasure does not always come with less pain; more goodness does not always come with less badness. So why think that anti-meritedness is special among all of the polar opposites in this regard?

Instead we should say that giving the medal to S_3 is not anti-merited, but we are tempted to think that it is because the act is unmerited partially in virtue of the presence of a feature (the cowardice of S_3 's conduct) that is the polar opposite of a feature that would, if obtained, partially make it the case that the act is merited: we confuse polar opposition among one of the grounds with polar

opposition of the property so grounded. Moreover, we can distinguish between the act of giving the medal to S_2 and the act of giving the medal to S_3 by saying that the latter act is more unmerited by a comfortable margin than the former; there is no need to say, in addition, that the latter act is “anti-merited,” for there is no such thing. Meritedness has only one opposite: unmeritedness.

6. Conclusion

Our specific topic has been the type of oppositeness relations born by the aptic properties, whether thick or thin. We have found that although some thick aptic properties—such as the desirable and the likeable—have polar opposites, most—such as the lovable and the persuasive—have only privative opposites. And we have found no polar opposites among the thin aptic properties, only privative opposites. The unfitting is the non-fitting, not the anti-fitting; the unmerited the non-merited, not the anti-merited; the inappropriate the non-appropriate, not the anti-appropriate. Unlike goodness, which has an inversely charged flipside, namely badness, aptic categories such as fittingness, meritedness, and appropriateness do not have inversely charged flipside, and we must take this characteristic into account when theorizing about them and in terms of them.

Our more general topic, though, has been the nature of opposition itself. Present-day philosophy’s silence about relations of opposition is puzzling. Relations such as identity and similarity are widely discussed and frequently invoked, so why not opposition? I hope to have shown that oppositeness is a serious topic about which we can theorize in a rigorous manner. I have put forward a number of hypotheses—for example, concerning the fundamental difference between privative, polar, and perhaps converse oppositeness—and made a number of assumptions—for example, concerning the relationship between privative opposites and relevant domains. But there are many questions about oppositeness that our discussion has raised without answering. Are there fundamental types of oppositeness beyond the ones broached here? Can anything substantive be said about what makes it the case that a category has an opposite of a given type? What determines the relevant domain for a given pair of opposites? Can one category have multiple opposites of the same type, or multiple opposites of multiple types? When and why do oppositeness relations transmit across certain definitional and explanatory structures but not others? Are there additional tests for determining types of oppositeness? Are there tests for determining whether two categories even are opposites, much less opposites of a given type? Is there a generalization of oppositeness in which a special kind of mutual incompatibility relation is born by three or more categories rather than by a pair of categories? To make progress on questions such as these, we would do well to look across

disciplines to the work contemporary linguists have done on lexical forms of oppositeness, as well as look back to earlier eras of philosophy when the nature of opposites was a central philosophical concern, as we find for instance in Aristotle's *Categories*, Plato's *Phaedo*, and the recorded thoughts of many of the pre-Socratics. But the relation between those works and the claims made here must wait for another occasion.

Acknowledgments

This essay grew out of pushback I received from philosophers when I took it to be obvious in (Berker 2022) that the unfitting is the non-fitting, the inapt the non-apt, the unmerited the non-merited, and so on. Many thanks to Arash Abizadeh, Claire Benn, Garrett Cullity, Nicholas Drake, Thomas Hurka, Conor McHugh, and Jonathan Way for questions and conversations that directly led to some of the material in this essay. Thanks as well to Rachel Achs, Jazlyn Cartaya, and Alex Friedman for written comments on an earlier draft; to audiences at the Tufts University, the University of Texas at Austin, and the 2022 Chapel Hill Metaethics Workshop for their remarks and queries; to the students in Elizabeth Harman and Sarah McGrath's Fall 2022 graduate seminar at Princeton University for their fantastic questions that led to a number of changes in the essay; and to Abdul Ansari, Robert Audi, Charles Barzun, Chris Blake-Turner, Britta Clark, Ian Cruise, Samuel Dishaw, James Doyle, James Dreier, Daniel Drucker, Matthew Evans, Caitlin Fitchett, Daniel Fogal, Alex Guerrero, Elizabeth Harman, Christopher Howard, Alex King, Harvey Lederman, Kathryn Lindeman, Sarah McGrath, Malcolm Morano, Jacob Mortimer, Daniel Muñoz, Mark Richard, Connie Rosati, Jacob Rosen, Becca Rothfeld, Stuart Shieber, Walter Sinnott-Armstrong, Matthew Vermaire, and Alex Worsnip for invaluable conversations on this material. Finally, I owe a debt of gratitude to two anonymous referees for *Ergo* and *Ergo* copyeditor Mark Steen for their thoughtful comments and suggestions.

References

- Achs, Rachel, and Oded Na'aman (2023). The Subtleties of Fit: Reassessing the Fit-Value Biconditionals. *Philosophical Studies*, 180(9), 2523–2546
- Anscombe, Gertrude E. M. (1957). *Intention* (2nd ed.). Harvard University Press.
- Berker, Selim (2022). The Deontic, the Evaluative, and the Fitting. In Christopher Howard and R. A. Rowland (Eds.), *Fittingness: Essays in the Philosophy of Normativity* (23–57). Oxford University Press.
- Bradford, Gwen (2021). Perfectionist Bads. *Philosophical Quarterly*, 71(3), 586–604.

- Bradford, Gwen (Ed.) (2022). Well-Being and Ill-Being. Special issue, *Midwest Studies in Philosophy*, 46.
- Brandt, Richard (1946). Moral Valuation. *Ethics*, 56(2), 106–121.
- Chappell, Richard Yetter (2012). Fittingness: The Sole Normative Primitive. *Philosophical Quarterly*, 62(249), 684–704.
- Cruse, D. Alan (1976). Three Classes of Antonym in English. *Lingua*, 38(3–4), 281–292.
- Cruse, D. Alan (1980). Antonyms and Gradable Complementaries. In Dieter Kastovsky (Ed.), *Perspektiven der lexikalischen Semantik* (14–25). Bouvier.
- Cruse, D. Alan (1986). *Lexical Semantics*. Cambridge University Press.
- Cruse, D. Alan (2002). Paradigmatic Relations of Exclusion and Opposition II: Reversivity. In D. Alan Cruse et al. (Eds.), *Lexicology: An International Handbook on the Nature and Structure of Words and Vocabularies* (Vol. 1, 507–510). Walter de Gruyter.
- D’Arms, Justin and Daniel Jacobson (2000a). The Moralistic Fallacy: On the “Appropriateness” of Emotions. *Philosophy and Phenomenological Research*, 61(1), 65–90.
- D’Arms, Justin and Daniel Jacobson (2000b). Sentiment and Value. *Ethics*, 110(4), 722–748.
- D’Arms, Justin and Daniel Jacobson (2023). *Rational Sentimentalism*. Oxford University Press.
- Diamond, Cora (2019). Reflections of a Dinosaur. *Proceedings and Addresses of the American Philosophical Association*, 93, 87–104.
- Gao, Chunming and Qianzhen Zheng (2014). A Linguistic Study of Antonymy in English Texts. *Journal of Language Teaching and Research*, 5(1), 234–238.
- Horn, Laurence R. (2020). Negation and Opposition: Contradiction and Contrariety in Logic and Language. In Viviane Déprez and M. Teresa Espinal (Eds.), *The Oxford Handbook of Negation* (7–25). Oxford University Press.
- Howard, Christopher (2019). The Fundamentality of Fit. *Oxford Studies in Metaethics*, 14, 216–236.
- Hurka, Thomas (2022). Against the Fundamentality of Fit. In Christopher Howard and R. A. Rowland (Eds.), *Fittingness: Essays in the Philosophy of Normativity* (58–79). Oxford University Press.
- Jones, Steven (2002). *Antonymy: A Corpus-Based Perspective*. Routledge.
- Joshi, Shrikant (2020). Affixal Negation. In Viviane Déprez and M. Teresa Espinal (Eds.), *The Oxford Handbook of Negation* (75–88). Oxford University Press.
- Kagan, Shelly (1998). *Normative Ethics*. Westview Press.
- Kagan, Shelly (2014). An Introduction to Ill-Being. *Oxford Studies in Normative Ethics*, 4, 261–288.
- Kjellmer, Göran (2005). Negated Adjectives in Modern English: A Corpus-Based Approach. *Studia Neophilologica*, 77(2), 156–170.
- Lee, Wooram (2022). Fittingness First?: Reasons to Withhold Belief. *Philosophical Studies*, 179(12), 3565–3581.
- Lehrer, Adrienne (1985). Markedness and Antonymy. *Journal of Linguistics*, 21(2), 397–429.
- Lehrer, Adrienne (2002). Paradigmatic Relations of Exclusion and Opposition I: Gradable Antonymy and Complementarity. In D. Alan Cruse et al. (Eds.), *Lexicology: An International Handbook on the Nature and Structure of Words and Vocabularies* (Vol. 1, 498–507). Walter de Gruyter.
- Lehrer, Adrienne and Keith Lehrer (1982). Antonymy. *Linguistics and Philosophy*, 5(4), 483–501.

- Lyons, John (1968). *Introduction to Theoretical Linguistics*. Cambridge University Press.
- Lyons, John (1977). *Semantics* (Vol. 1). Cambridge University Press.
- Maguire, Barry (2018). There Are No Reasons for Affective Attitudes. *Mind*, 127(507), 779–805.
- McHugh, Conor and Jonathan Way (2016). Fittingness First. *Ethics*, 126(3), 575–606.
- McHugh, Conor and Jonathan Way (2022). *Getting Things Right: Fittingness, Value, and Reasons*. Oxford University Press.
- Murphy, M. Lynne (2003). *Semantic Relations and the Lexicon*. Cambridge University Press.
- Nagel, Thomas (1970). *The Possibility of Altruism*. Princeton University Press.
- Palmer, Frank R. (1976). *Semantics: A New Outline*. Cambridge University Press.
- Quirk, Randolph, et al. (1985). *A Comprehensive Grammar of the English Language*. Longman.
- Rice, Christopher M. (2019). Objective List Theories and Ill-Being. *Ethical Theory and Moral Practice*, 22(5), 1073–1085.
- Rosa, Luis (2023). Suspending Judgment the Correct Way. *Inquiry*, 66(10), 2001–2023.
- Srinivasan, Amia (2020, August 24). What Have We Done to the Whale? *New Yorker*. Retrieved from <https://www.newyorker.com/magazine/2020/08/24/what-have-we-done-to-the-whale>
- Sumner, Wayne (2020). The Worst Things in Life. *Grazer Philosophische Studien*, 97(3), 419–432.
- Vollmer, Michael (forthcoming). Suspension of Judgment: Fittingness, Reasons, and Permissivism. *Episteme*.
- Zimmer, Karl (1964). *Affixal Negation in English and Other Languages*. Supplement to *Word*, 20(2) (Monograph 5).
- Zimmerman, Michael J. (2002). Taking Luck Seriously. *Journal of Philosophy*, 99(11), 553–576.
- Zimmerman, Michael J. (2006). Moral Luck: A Partial Map. *Canadian Journal of Philosophy*, 36(4), 585–608.