

Reasons: Wrong, Right, Normative, Fundamental^{*†}

Errol Lord

University of Pennsylvania

Kurt Sylvan

University of Southampton

1 Introduction

According to *reasons fundamentalism*, all normative properties are analyzable in terms of reasons.¹ Famously, some of the analyses offered by reasons fundamentalists face the *wrong kind of reasons problem*. This problem first appeared in the literature on the buck-passing account of value, which says in its simplest form that what it is for something to be valuable is for there to be sufficient reasons to have a pro-attitude toward it.² This simple view fails, many worry, because there can be reasons for having pro-attitudes towards things that have nothing to do with their value. Contrasting cases like Beauty and Extra Credit provide an illustration:

Beauty: Jane is a first-year graduate student in Art History. She has loved art all her life, but is just now getting the opportunity to see Europe’s masterpieces through her graduate program. She sees the Mona Lisa in person for the first time. She is enthralled by its symmetry, depth, and enigmatic tone.

Extra Credit: Jack is a smart aleck on a high school trip to Paris. He couldn’t care less about art and finds the Louvre to be a total bore. Predictably, he has been causing all

^{*}forthcoming in *Journal of Ethics & Social Philosophy*

[†]Thanks to audiences at the Southern Normativity Group, SLACRR, and the a workshop hosted by the Varieties of Normativity Project, especially Matti Eklund, Jack Woods, Stephanie Leary, Pekka Väyrynen, Daniel Fogal, Kathryn Lindeman, Amelia Hicks... Thanks also to Nate Sharadin and Barry Maguire for crucial inspiration at the beginning.

¹Prominent reasons fundamentalists include Scanlon (2014), Parfit (2011), Schroeder (2007), Skorupski (2011) and Dancy (2004). We are here understanding reasons fundamentalism as a claim about the internal structure of the normative—i.e., as the claim that reasons are the fundamental items *within the normative domain*. It is consistent with reasons fundamentalism of the sort considered here that reasons might in turn out to be naturalistically grounded. While “reasons fundamentalism” could be used in a stronger sense that implies non-naturalism, that is not how we are using it.

²See Rabinowicz & Rønnow-Rasmussen (2004), D’Arms & Jacobson (2000a,b).

sorts of trouble. As they enter the room with the Mona Lisa, Jack starts photo bombing tourists' shots of the painting. In a fit of frustration, Jack's high school history teacher says to him, "Will you just be enthralled by *something*? I will give you extra credit if you show some appreciation and respect for the Mona Lisa."

Jane responds to some reasons to be enthralled by the Mona Lisa. Intuitively, her reasons are of the right kind. The painting's depth, symmetry, and enigmatic tone are part of what make it valuable. Since Jane's enthralment is a response to these reasons, it is plausibly *fitting*. Jack's reasons are different. The fact that Jack will get extra credit if he is enthralled does not help to explain why the Mona Lisa is valuable. Jack's teacher could have provided a similar reason for him to be enthralled by something with no value—a saucer of mud, for example. If Jack were to be enthralled just to get extra credit, it is plausible that his enthralment would be unfitting. This is because he is enthralled for the wrong kind of reason.

Buck-passers about value need to understand value solely in terms of the right kind of reasons, it seems.³ Otherwise it seems their view badly overgeneralizes. A sufficient incentive to have a pro-attitude would entail that the object of that pro-attitude is valuable. This is a bad prediction. So buck-passers face pressure to explain the distinction between the right kind of reasons (RKR) and the wrong kind of reasons (WKR).

While the literature on buck-passing is the most familiar context in which the wrong kind of reasons problem arises, it is easy to see that the same problem arises for other reasons fundamentalist analyses. To see why, one need only consider how reasons fundamentalists might go about analyzing other normative properties that can be picked out with words of the form ' ϕ -ible/-able' and ' ϕ -worthy', like the properties of being *credible*, *trustworthy*, *despicable*, *noteworthy*, and so on. In each case, it is plausible that there can be sufficient reasons for ϕ -ing an X that don't bear on X 's ϕ -ability. For this reason, it is clear that the wrong kind of reasons problem is a much more general problem.

Interestingly, this point shows not only that reasons fundamentalists have many WKR problems, but also that *many theorists* have WKR problems. For many theorists can and should find it attractive to analyse *some* things in terms of reasons. Indeed, as Schroeder (2010) emphasizes, it is attractive to think that a reasons-based account has to be the right account of properties like credibility, despicability, trustworthiness, noteworthiness, and so on, whatever one thinks about the broader reasons fundamentalist program. Reflections on this fact have led some—most notably, Schroeder (2010)—to suggest that fundamentalists are off the hook. The fact that their account would badly overgeneralize without a principled distinction between RKR and WKR does not undermine their approach, they insist, because the WKR problem is everybody's problem, and the challenge for fundamentalists will be answered if the more general question has an answer. And so they conclude that they face no special challenge, thanks to the generality of the problem.

³We add "it seems" because some buck-passers (e.g., Skorupski (2011), Parfit (2011)) deny that there are wrong-kind reasons, and hence appeal to reasons, period, rather than right-kind reasons. But this response neglects the generality of the phenomenon, as we will see.

Our main goal in this paper is to argue that this conclusion is mistaken and indeed gets things backwards. We think that reflection on the alleged generality of the distinction between right- and wrong-kind reasons suggests that life is *harder*, not easier, for reasons fundamentalists. With this goal in mind, we defend two main claims. The first claim is:

(1) The apparent generality of the WKR problem suggests that the class of right-kind reasons is wider than the class of normative reasons; the reasons that determine “fittingness” are not essentially normative in any sense stronger than that associated with any arbitrary standard of correctness. This gives rise to what we call the *Right-Kind of Reasons Problem*.

The second claim is:

(2) If (1) is true, then some special problems—including a potentially fatal dilemma—arise for fundamentalists.

After defending (2), we will evaluate several possible escape routes for fundamentalists.

Our plan is as follows. In §2, we begin by considering the case for the generality of the WKR problem at greater length. We argue that any adequate solution to the more general WKR problem will imply that not all RKR are normative reasons. Thus, we cannot determine which considerations properly figure into reasons-based analyses simply by solving the WKR problem. We also need to figure out which RKRs are normative reasons, and why they are normative reasons. This is the Right-Kind of Reasons Problem. To solve the RKR problem, one needs to solve two other problems, which we call the *Normativity Problem* and the *Rationale Problem*. In §3, we argue that reflection on these problems raises a dilemma, which we call the *Fundamentalists’ Dilemma*. In §4, we canvass three responses to this dilemma. The first two are reductive naturalist views—the first, a Humean view, and the second, a novel constitutivist view. The third is a new form of quietism, one which enjoys some advantages over other quietist views. We suggest that the constitutivist option is the best. We leave it to future work to decide whether this is a strike against reasons fundamentalism or an argument for constitutivism. In §5, we summarize the key morals and conclude.

2 Why the Generality of the RKR/WKR Distinction Makes Life Harder for Fundamentalists

2.1 The Generality of the Distinction and Why Not All RKRs are Normative Reasons

If the WKR problem is everybody’s problem, there must be a distinction between RKRs and WKRs that all theorists can and should acknowledge. Thankfully for reasons fundamentalists, there is a plausible candidate for such a distinction, and it is a *very* general distinction. As

Schroeder (2010) points out, wherever one finds a *standard of correctness* governing an activity, one finds a distinction among considerations recommending potential moves in the activity that could naturally be dubbed a distinction between reasons of the right kind and reasons of the wrong kind.⁴ In particular, for any standard of correctness, there are certain considerations that bear on whether an act would help one to conform to the standard and others that are simply irrelevant. Insofar, then, as one is engaged in an activity governed by a standard of correctness, there are certain considerations to which one ought to be sensitive relative to the standard, and others to which one oughtn't be sensitive relative to the standard.

It is natural enough to describe this distinction as a distinction between reasons of the right kind and reasons of the wrong kind. To appreciate this fact, one need only consider examples like the following pairs of cases:

Rope Pair

Good Boy Scout: Kenny is a precocious and studious boy scout. He has learned how to tie most of the knots in the Boy Scout handbook. He is currently trying to tie a half-hitch. Placing the left portion of the rope over the right would be an efficient step toward producing a half-hitch as described by the book. Kenny chooses to manipulate the rope accordingly with this fact in mind.

Bad Boy Scout: Billy is a terrible boy scout. He has it out for Kenny, and likes to mess with Kenny's sense of Boy Scout decency. So when he sees Kenny practicing his knot tying skills, he decides to have some fun with him. He decides to offer him \$20 to deviate from the book when tying his half-hitch. Kenny sees the utility of the \$20 and is thus disturbed, just as Billy intended. Kenny gives in this one time and moves the rope in a way that will at best lead to a very bad example of a half-hitch.

Table Pair

Proper Butler: Mr. Carson is a proper English butler. He always sets the table to the standards of the proper English tradition. Arranging the forks to the left of the knives now is one way to help get the table in order. He chooses to arrange the silverware accordingly with this fact in mind.

Devious Valet: Thomas is a devious valet. He doesn't care much for the traditions except when he feels they are to his advantage. He decides one day that it is to his advantage to incentivize breaking tradition for Mr. Carson. So he makes it so that the only way for Mr. Carson to avoid embarrassment to Lady Mary is by breaking with table setting tradition and putting the knives to the left of the forks.

⁴This point is central to both Schroeder's (see especially Schroeder (2010)) and D'Arms and Jacobson's (see especially D'Arms & Jacobson (2000b)) approaches to the WKR problem.

In Good Boy Scout, Kenny's choice is appropriate relative to the standard of correctness established by the handbook's description of a half-hitch (though he hasn't at the end of the example yet produced a half-hitch).⁵ Why? The natural explanation is that Kenny's choice properly takes into account how various available ways of manipulating the rope would make a difference to whether he correctly produces a half-hitch. Considerations that bear on whether a certain way of manipulating the rope would correctly produce a half-hitch are reasons of the right kind relative to the standard. Hence, Kenny's choice is appropriate relative to the standard because it is sensitive to reasons of the right kind relative to the standard. In Bad Boy Scout, Kenny is not doing what he should be doing by the lights of the standard. He is responding to an irrelevant consideration relative to the standard, by choosing to move the rope in a way that will merely produce some money for him.

Similar things can be said about the Table Pair. In Proper Butler, Mr. Carson's choice is appropriate by the lights of the relevant standards. A natural explanation is that he chooses for reasons that are relevant and good relative to the standard: he chooses to move the silverware in a way that will help him to set the whole table in the proper English way. In Devious Valet, Mr. Carson is led off the proper path by considerations that are irrelevant to the norms governing such arrangements.

These cases draw attention to a distinction that seems worth calling a distinction between reasons of the right kind and reasons of the wrong kind. Notice that this distinction is not just a distinction among motivating or explanatory reasons. Before Kenny decides to place the left portion over the right, there are various moves open to him. In choosing which move to make, there are facts to be borne in mind, like:

(RR) The fact that placing the left portion over the right is the most efficient way to complete the first step towards correctly tying a half-hitch.

To those attempting a half-hitch, RR recommends placing the left portion over the right at Kenny's stage. If it were causally possible to skip this step, RR might not conclusively recommend the move. Accordingly, RR's apparent force may vary depending on the causal possibilities. But assuming there is no way to skip the step, RR does conclusively recommend the move to the half-hitcher. And it does so even if Kenny hasn't yet registered this fact, and hence isn't among his motivating reasons.

To the extent, then, that it is plausible to call this fact a reason of the right kind, it seems not to be merely a motivating or explanatory reason. It is a consideration carrying a recommending weight that modulates the strength of the criticism the half-hitcher would face if he

⁵Note that we here assume that a choice or act can be *appropriate* relative to a standard of correctness even if it doesn't constitute *conformity* to the standard. We will also use the word "fitting" in this way, though we realize that some use "fitting" as a synonym for "correct". We assume that appropriateness can be naturally enough analyzed in terms of reasons of the right kind relative to the standard, where these are understood as considerations that bear on whether acting in the relevant way would make a contribution toward conformity. We do not similarly assume that the standard of correctness is analyzable in terms of reasons; indeed, we will ultimately suggest that as long as standards of correctness are not inherently normative, it is permissible for a reasons fundamentalist to explain the right-kind/wrong-kind reason distinction by appeal to them.

pursued a more cumbersome strategy. And the criticism needn't be that he doesn't conform to the standard: he might eventually bumble into conformity. This fact suggests that the right-kind/wrong-kind distinction at issue isn't just a distinction between whether or not one is doing what conforms to the standard. Each choice and move can be assessed for a kind of appropriateness or fittingness relative to the standard. This appropriateness is a function of how well one is responding to considerations relevant to the attainment of the standard.

Observations of this kind make it plausible not just that there is an RKR/WKR distinction that everyone can and should acknowledge, but also that there is a task worthy of everyone's attention: the task of explaining why certain considerations are relevant to this distinctive kind of criticism, while others are irrelevant despite having great, even conclusive, normative significance. Indeed, this task is compulsory for anyone interested in the grounds of the criticism. Again, it is plausible that whether a move would be the most appropriate move relative to the standard is determined by whether it is the move most strongly recommended by these considerations, where some considerations can recommend with less than sufficient force (e.g., *this move would help a bit (though others are more helpful)*).

Of course, this task doesn't look daunting: the difference is clearly tied to the bearing of the considerations on whether one would meet the standards that constitutively govern the activity. Precisely this fact, one might hope, saves reasons fundamentalists from their problem. For, one might hope, their problem can then be solved as a special case of this easier and more general problem.

Appreciating this point, Schroeder (2010) uses the generality of the problem as inspiration for a particular account of what it is for something to be a right kind of reason. According to this account, the right kind of reasons are always relativized to *activities*. They are the reasons that everyone engaged in the activity has *in virtue of* engaging in that activity. More officially:

Activity: What it is for r to be a right kind of reason to ϕ relative to some activity a is for r to be a reason for anyone engaged in a to ϕ in virtue of being engaged in a .

Activity is poised to make plausible extensional predictions. It doesn't seem like *anyone* tying knots has the reason provided by Billy—only Kenny does. But anyone tying Boy Scout knots does have the reasons provided by the book. The nature of Boy Scout rope tying seems to guarantee this. Similarly, not all the good butlers have the reason provided by Thomas—only Mr. Carson does. But all butlers have the reasons provided by tradition. The nature of butlering seems to guarantee this.

There are two basic thoughts behind Activity. First, wrong-kind reasons aren't tied to the nature of the activities in the right kind of way. Second, there is something too local about wrong-kind reasons. Extensional predictions aside, we think that there is something very powerful behind these thoughts. The wrong kind of reasons don't essentially bear on whether one is helping to fulfill the aim of the activity, while the right kind of reasons do, and are considerations that everyone engaged in the activity ought to consider on pain of meriting a kind of criticism. Although we regard these thoughts as genuine insights, we will not assume that

Activity in particular is true.⁶ We merely assume in what follows that some account that vindicates these insights has to be right, and Activity is at least a good stab in the right direction.

Before proceeding further it is worth considering an objection to Activity in order to highlight what we will be assuming. The objection is that Activity puts the cart before the horse by insisting that what is fundamental are the activities; what is fundamental, according to the objector, are the reasons themselves. We can then derive activity-types from the reasons however we wish. For example, the objector might insist that what is fundamental to epistemic normativity are reasons for certain reactions. The activity of believing will then just be whatever activity purports to be sensitive to such reasons.

Reflection on standard buck-passing accounts of value encourages this objection. For such accounts appear to proceed on the assumption that what is fundamental are the reasons for various reactions. Once those reasons are picked out, we can then determine the activities (if we are so inclined). Activity denies that this is the order of explanation. Instead, it holds that the activities come before the (right-kind) reasons. The nature of the activities themselves determine the right-kind reasons.

While we recognize that this picture is natural—especially upon reflection on the basic structure of buck-passing accounts of value—we think that it ultimately misses the key lessons of the WKR problem. The fundamental lesson of the WKR problem is that *something* must be said in order to pick out the facts that are the reasons that serve in analyzing properties like value. This lesson of course does not force us to Activity. But once this point is appreciated, the virtues of Activity become apparent, for it does seem plausible that there is a *correlation* between certain activities and the right-kind reasons. When you couple this fact with the need to explain why right-kind reasons are right-kind, it starts to look appealing to think that the nature of the activities are what explain why the right-kind reasons are right-kind.

That said, we are not endorsing Activity. We are only endorsing four commitments that Activity helps to illustrate: first, that we need an explanation of what makes right-kind reasons right-kind; second, that one can draw the right-kind/wrong-kind distinction wherever there are standards of correctness; third, that the WKRs are wrong-kind in virtue of not being tied to the standards in the right way; fourth, and finally, that the RKR's aren't local and idiosyncratic: rather, they are reasons shared by those who are tied to the standards in some relevant way.⁷ Since Activity illustrates the last three commitments so nicely, we will largely proceed as if it is true. This is merely for the sake of concreteness and simplicity. Assuming a concrete view

⁶We will also not be assuming Schroeder's account of what it takes for r to be a reason for anyone engaged in a in virtue of being engaged in a . In fact, we think that Sharadin (2013, 2016) shows that Schroeder's view cannot work. We, like Sharadin, think that this sort of agent-neutrality should be explained in terms of the constitutive standards of the activity. We suspect Schroeder doesn't go for this because he is afraid that this will be smuggling normative notions into his account of RKR's, which would in turn undermine his fundamentalism. We think that this is a mistake because it is a mistake to think that constitutive standards are genuinely normative. This will be spelled out in more detail below (see especially §4.2).

⁷We are fine with one reading our conclusion in this section conditionally. *If* one accepts these commitments, *then* it is plausible that not all right-kind reasons are normative reasons. Establishing this conditional would still be an important result given that prominent views accept the commitments.

will make it easier to state the problems at the heart of the paper.

Given these commitments, however, we must conclude that not all RKR are normative reasons. For standards of correctness *per se* are normative only in an anaemic sense. They set up some standard, but not all are genuinely normative. Explaining the distinction between anaemic and genuine normativity is a hard task, but there are clear examples on either side: norms of etiquette and grammar, on the one hand, and epistemic norms, prudential norms, and moral norms, on the other.⁸ All we assume here is that being a standard of correctness isn't sufficient for being genuinely normative. This claim is one that even fundamentalists who think normativity comes fairly cheaply can embrace: for no fundamentalists would think there is necessarily a good reason to satisfy any standard whatsoever (though some believe in a vaster array of reasons than others). Yet it seems all standards generate a distinction between right-kind and wrong-kind reasons. So, for the standards that aren't genuinely normative, we get RKR that are not normative reasons.

To see this point vividly, consider cooking. There are many standards of correctness that purport to govern the cooking of *cacio e pepe* (a Roman pasta dish—literally, 'cheese and pepper').⁹ For there are loads of recipes for making this dish. Many of them conflict. For any given standard corresponding to a recipe, there will be RKR to do certain things—e.g., to add extra cheese table-side. But—and here is the important point—not all of these RKR are normative reasons, even for those who are engaged in the cooking of *cacio e pepe*. As any self-respecting Roman will tell you, many of these recipes lead to inferior *cacio e pepe*. The RKR associated with those standards plausibly have no more genuinely normative significance than other anaemic standards, such as the norms of etiquette and grammar.

It does not follow from this point that the right-kind/wrong-kind distinction that all should acknowledge is not a distinction in *reasons*, as we will again emphasize in §2.4. There is a criticism to which one is open in ignoring considerations relevant to satisfying the standards of correctness governing an activity. And these considerations have weights linked to their bearing on how well one would be satisfying the standards that govern the activities, in the attributive-value sense of "well". They hence have the functional profile of reasons: they are recommending considerations with weights whose balance determines a kind of criticism. But the genuine normative significance of both the criticism and the considerations that help to explain it is as much an open question as the normativity of criticisms of instrumental rationality and of the subjective reasons associated with requirements of instrumental rationality (which we also assume are indeed reasons, though perhaps not always normative reasons if Humeanism is false!).

Indeed, there is pressure for the reasons fundamentalist who wants to write off the WKR problem as everybody's problem to insist that this distinction is a distinction among reasons. For the mere fact that a more general wrong-kind/right-kind distinction can be drawn doesn't

⁸Perhaps the most famous invocation of the difference between genuine normativity and anaemic normativity is Foot (1972). For further discussion, see Wodak (FC), Parfit (2011), Copp (2004), Woods (FC), Baker (FC), Plunkett & Shapiro (2017), McPherson (2011).

⁹In EL's opinion, the most delightful thing one can put in one's mouth when it is prepared well.

show that there is a problem that everyone faces parallel to the WKR problem. A genuine analogue of the problem would arise only if there were a kind of assessment which ought to be understood as determined by the balance of considerations worth calling “reasons”. For only if there is independent reason for everyone to analyze some form of assessment in terms of reasons of the right kind would it be true that there is a WKR *problem* for everyone. It is not enough that there be a distinction that everyone can understand. For there might be a special obligation for certain theorists to explain this distinction owing to their need to invoke it in their theory, and only this fact generates a problem.

The upshot is that the generality of the WKR problem shows that not all RKR are normative reasons, for some standards of correctness are normatively anaemic. These standards generate RKR that are not plausibly normative reasons. So, in order to determine which considerations are relevant for reasons-based analyses—in order to figure out what the normative reasons are—one has to determine which RKR are normative reasons, and why they are normative reasons. This is the *Right-Kind of Reasons problem*. It has the same form as the WKR problem, for that too is a problem concerned with differentiating some favorers from others. It is widely assumed that once we sort the wrong-kind favorers from the right-kind favorers, we will know which favorers to appeal to in reasons-based analyses. We’ve argued that this is false; we also need to know which right-kind favorers flow from robustly normative standards. To provide this story is to provide a solution to the RKR problem.

2.2 How to Solve the Right-Kind of Reasons Problem

We agree that the WKR problem is everybody’s problem. For similar reasons, we think that the RKR problem is everybody’s problem.¹⁰ In this subsection we will provide a schema for solving it.

Recall how we got stuck with the WKR problem. We start trying to analyze complex normative properties in terms of normative reasons. We then see that there are some facts that seem to recommend reactions in the wrong way. This observation leads us to search for an account of what it is for a fact to recommend a reaction in the right way. But as it turns out, a fact can recommend in the right way and yet *still* not be a normative reason. So, solving the WKR problem isn’t sufficient for being in a position to pick out the normative reasons that are supposed to be analyzing the complex stuff. This extra problem is the RKR problem.

To solve the RKR problem we need to carry out two extra steps. First, not only do we need an account like Activity, we also need an account about which activities are normatively relevant.¹¹ That is, we need to know which activities, if any, are such that the RKR associated with those activities are necessarily normative reasons. Once we have an account of this, will

¹⁰As we mentioned above, we are happy to weaken this to the claim that the RKR problem is a problem for everyone who accepts that the right view of RKR has the four features we elucidate in §2.1.

¹¹Again, we will speak of activities because we find it illuminating. One needn’t appeal to activities per se. One could also appeal to standards of correctness or perhaps some other notion that plays the role that activities play in Activity.

we be in a position to determine which favorers can figure into reasons-based analyses. We call the problem of determining which activities are robustly normative the *Normativity Problem*. Solving this problem is the first step towards solving the RKR problem.

In order to solve the Normativity Problem, one must have a story about which activities are normatively relevant. Now, we recognize that there are some obvious candidates: believing, intending, and desiring, for example, will plausibly be on the list. The hard part is not coming up with a *list* of activities that are normatively relevant. So solving the Normativity Problem doesn't look that daunting. However, it is harder to give an account of what these activities have in common that makes them normatively relevant. We call this the *Rationale Problem*. Solving the Rationale Problem is the second step towards solving the RKR problem. With a solution to both the Normativity and the Rationale problems, one will have an account of which activities are robustly normative and an explanation of why *those* are the robustly normative activities.

So far problems have only multiplied. We started off with the WKR problem. Now we have the RKR problem, which naturally breaks down into the Normativity Problem and the Rationale Problem. As we've said, the WKR and RKR problem are structurally on all fours. They are problems for everyone interested in using normative reasons to analyze normative phenomena. Further, fundamentalists face both problems in full generality. In other words, the success of the fundamentalist's entire metanormative theory hinges on solving both problems. Without a solution to both problems, the fundamentalist lacks the resources to even get going. This makes it all the more unfortunate that the fundamentalist has a particularly hard time solving the RKR problem. We'll see this shortly. First, though, we will elucidate the RKR problem more by considering an analogous problem that only targets buck-passing about value. This will naturally lead to an attempt to dismiss the problem out of hand. After dismissing that dismissive response, we will explain why the fundamentalist has a particularly hard time solving the RKR problem.

2.3 Analogous Problems for Buck-Passers about Value *Simpliciter*

To further understand the RKR problem, it is helpful to compare it to an analogous but more local problem for a particular fundamentalist analysis—namely, the buck-passing account of value *simpliciter*.¹² In examining this more local problem, we will discover some points that help to address a dismissive response to the RKR problem, which we rebut in the next subsection.

According to the simple buck-passing account of value *simpliciter*, what it is for *X* to be valuable *simpliciter* is for there to be sufficient reasons to have some pro-attitude toward *X*. This simple account, of course, faces the WKR problem. But as Schroeder points out, it is

¹²We don't assume that reasons fundamentalists are committed to the existence of value *simpliciter*, or that buck-passing accounts of value are exhausted by buck-passing accounts of value *simpliciter*. We focus on this case to explain the kind of problem we are raising, which could also be raised for buck-passing accounts of goodness-for (though perhaps not for accounts of attributive goodness, which isn't clearly normative).

plausible that everyone should accept a buck-passing account of certain evaluative properties, like enviability and amusingness. He concludes that for this reason the buck-passer is off the hook.

Although considering properties like admirability and enviability points to a WKR/RKR distinction that everyone should want to explain, life is harder for buck-passers about value *simpliciter* for this very reason. The problem cases that arise for these accounts point to a notion of a right-kind reason and a corresponding notion of fittingness unsuited for underpinning a buck-passing account of value *simpliciter*. This suggests that the real problem for buck-passers—that of explaining why certain reasons are relevant to the grounding of value *simpliciter* and others aren't—remains even after they produce an adequate account of the WKR/RKR distinction. This is the localized version of the RKR problem.

To see this clearly, consider first the example of 'amusing'. Close cousins to famous points made by Justin D'Arms and Daniel Jacobson bring out the challenge.¹³ The key point is that the notion of a right-kind reason that makes the following analysis true is insensitive to goodness and badness *simpliciter*:

Amusing: What it is for x to be amusing is for there to be right-kind reasons to be amused by x .

The features of a joke or situation that make it amusing may well be features that make it bad *simpliciter*. There is something funny about someone's repeatedly making a prudentially bad kind of mistake (this is the whole idea of slapstick comedy), and about absurd and incongruous instantiations of badness (imagine a house being struck by lightning 100 times in a row). Fans of dark humor will find the absurd badness of the world funny in itself. On a natural first reading of the buck-passing account, this seems impossible. For the buck-passing account holds that the features we have reason to have pro-attitudes like amusement towards are good *simpliciter* (and in virtue of the fact that we have such reasons). This appears to ignore the role that the internal standards of funniness play. In particular, it seems to ignore that those standards seem insensitive to goodness and badness *simpliciter*. Those standards only care about humor, as it were. They embrace it wherever it is found, even in the bad.

One could try to deny that these things are amusing. But to do so would, as D'Arms and Jacobson stress, involve pointing to features that are wrong-kind reasons relative to the standards that spell out what it takes for something to be funny. That is, to think that these things are not funny simply because they are bad *simpliciter* is akin to thinking something is not funny because it is morally bad. It is to point to wrong-kind reasons relative to the standards governing funniness. Our point, again, is that the standards governing funniness are not directly sensitive to goodness and badness *simpliciter*. This is a cousin of D'Arms and Jacobson's point that the standards of funniness are not directly sensitive to moral badness/impermissibility.¹⁴

¹³See D'Arms & Jacobson (2000a).

¹⁴It is worth spelling out the exact relationship our point has to D'Arms and Jacobson's main point. Their main point is that it is a conceptual mistake to think that the fact that something is morally bad/impermissible

Importantly, this is not a local point about amusement. It is plausible that emotions present their objects as meeting or failing to meet some standard of correctness, in some sense of “present”. But it is easy to imagine emotions that present their objects as meeting standards of correctness that fail to line up with anything good or bad *simpliciter*. More familiar emotions provide further examples. Consider emotions like *shock* at x and their corresponding correctness conditions (i.e., x 's being *shocking*). Or if you don't like that example, consider religious emotions like *penitence*, which is fitting only if one has sinned. There are conditions under which these emotions are fitting, but from this nothing follows about whether anything of normative significance is encoded by these feelings.

It is easy to imagine other possible emotions that would illustrate the same moral. We might not have names in ordinary language for some of these emotions. But we may have other signs—consider an emoticon with someone wearing sunglasses giving a thumbs-up, which represents that to which it is responding as *cool*. One could imagine parallel emoticons for *fly*, *swag*, *rad*, etc. Supposing that we could coin words for these emotions of the form EMOT, we could imagine corresponding adjectives of the form EMOT-able. Right-kind reasons for EMOT-ing X will be evidence that X is EMOT-able. But the EMOT-ability conditions will be given by a standard of correctness that may not map onto anything of value or disvalue *simpliciter*.

What these reflections suggest is that we need a distinction between right-kind and wrong-kind considerations in favour of pro-attitudes and con-attitudes that flows from their built-in standards, without regard to whether these standards track value *simpliciter*. But if so, then the following revised version of the buck-passing account is false:

Right-Kind Buck-Passing about Value *Simpliciter*: x is valuable (/disvaluable) *simpliciter* iff there are sufficient right-kind reasons for having some pro-attitude (/con-attitude) toward x .

We can see that this cannot be right by thinking about cases like amusement, envyability, and shock. Amusingness, envyability, and shockingness do correlate with right-kind reasons for the attitudes of amusement, envy, and shock. But there is an open question about whether these properties are values/disvalues *simpliciter*. If so, we should reject Schroeder's case for getting buck-passers about value *simpliciter* off the hook. Perhaps everyone should expect there to be a solution to the WKR problem for ‘amusing’ and ‘shocking’. But it won't give the buck-passer everything she needs. It actually shows that she needs more than the general distinction between WKR/RKR for pro-attitudes. She also needs to tell us which pro-attitudes are relevant to the analysis of value *simpliciter*.

is a reason not to be amused. We are making a parallel point against a reply the buck-passer might give to our suggestion that features that are bad *simpliciter* cannot be funny. To give that response is to ignore the fact that the standards that govern what is funny do not analytically rule out that features that are bad *simpliciter* are funny in virtue of being bad *simpliciter*. We are also making a point that goes beyond D'Arms and Jacobson's, namely that paying attention to the internal standards of amusement puts pressure on the buck-passer to explain which pro-attitudes count. This is because our point about amusement shows that amusement is not a pro-attitude that should ground goodness *simpliciter*. This illustrates the buck-passers analogue to the Normativity Problem by showing that she has the task of telling us which pro-attitudes count.

The buck-passer about value *simpliciter* now faces two problems. First, she must tell us which pro-attitudes are such that right-kind reasons for them ground value *simpliciter*. This is her analogue of the Normativity Problem. Second, the buck-passer must explain what these attitudes have in common, such that the presence of sufficient RKR for bearing them to x entail that x valuable. This is the buck-passer's analogue of the Rationale Problem.

Notice that the first problem differs from a superficially similar "problem" sometimes mentioned for buck-passers. It has, of course, always been a *question* for buck-passers which pro-attitudes are the ones to which they should at bottom appeal. Sometimes this question is regarded as raising a problem. But there is not a problem—just an interesting question—until it is shown that not all right-kind reasons for pro-attitudes are value-grounding in the robust sense of "value" that buck-passers were originally trying to analyze. But we can see from reflecting on the generality of the WKR problem that not all right-kind reasons for pro-attitudes are value-grounding in that sense. Hence, it is incumbent on buck-passers to solve their analogues of the Normativity and Rationale Problems.¹⁵

2.4 Dismissing a Dismissive Response

Reflection on the foregoing special case helps to address a dismissive response to our problems. One might have imagined some reasons fundamentalists, such as Parfit or Scanlon, agreeing that there is a right-kind/wrong-kind distinction that arises relative to any arbitrary standard of correctness, but insisting that this distinction is not necessarily a distinction *among reasons*. The right-kind considerations that we are claiming are not normative reasons are, these theorists might claim, *not reasons at all*. So, these theorists might insist that they don't have to say which attitudes or activities are relevant. Call this the *dismissive response* to the RKR problem.

Before giving our main response to this reply, it is worth noting that this reply is not available to all fundamentalists. Primitivist fundamentalists like Parfit and Scanlon can make this move, but fundamentalists like Schroeder who are also Humeans cannot. Schroeder does think—unsurprisingly, given his Humeanism—that one can have reasons to fulfill the aim of some activity simply in virtue of being engaged in that activity. If one is trying to tie a knot, one presumably wants the knot to be tied and that desire just does, on a Humean view, generate certain reasons. (We will return to this below.) While Parfit and Scanlon might deride activity-relative reasons, such derision would be theory-driven and, we think, not terribly convincing.

More importantly, the Dismissive Reply is implausible as a response to the localized

¹⁵One option, of course, is to insist that the pro- and con-attitudes that play the relevant theoretical roles in the buck-passing account are primitive. This would side-step the local versions of the Normativity and Rationale objections raised here. We have two things to say about this. First, this wouldn't show that the problems aren't problems for the buck-passer. Indeed, to accept this form of primitivism is a way of accepting that there is a problem. Second, primitivism is the option of last resort. We should investigate whether there are viable explanations before we become primitivists. We think that there are views that offer explanations. We explore some below.

problem for the buck-passing account. Some reasons fundamentalists might antecedently have been fine denying that right-kind considerations in an activity governed by some standard of correctness are right-kind reasons in that activity. But it is not plausible to deny that the amusing properties of a joke are reasons to be amused by it. It makes good sense to be amused by things that are amusing, whatever else might be said about one's amusement. The properties of a joke can also provide more or less of a case for being amused, and hence play a contributory favouring role, as reasons do. We see no theory-neutral reason to deny the strong intuition that if a joke is amusing, there are reasons to be amused by it. Since features that are bad *simpliciter* aren't necessarily right-kind reasons not to be amused (and some things may be amusing *because they are absurdly bad*), the funny-making features will not be relevantly defeated by other right-kind reasons. So, the point remains there: sufficient right-kind reasons for some pro-attitudes do not generate value *simpliciter*.

It is worth emphasizing that the point we are making here is a more local one that is intended to clarify the structure of our main point. Although it is structurally parallel to our main point, it is not the same point. What we are arguing here is that the existence of sufficient reasons of the right kind for having a pro-attitude toward x doesn't imply that x is valuable *simpliciter*. Even if one were to agree that the reasons in these cases are genuinely normative, our conclusion here would remain. It is unclear that being amusing entails being good *simpliciter* to any degree or in any way, or that being enviable entails being good *simpliciter* to any degree or in any way. We don't think it is plausible to claim that a deeply unequal world would be good in a way because it would entail that some people are enviable, where this is a way of being good. (For this reason, the analogue of our Normativity Problem presently under discussion might better be called the Value Problem.)

Accordingly, the problem here is to explain which right-kind reasons for pro-attitudes are relevantly value-grounding, given that not all are. We think this problem is illuminatingly analogous to the problem of explaining which right-kind reasons in general are relevantly normativity-grounding. But the main difference between the problems simply reinforces the strength of the original Normativity and Rationale Problems. For the assumptions needed to get these problems going are much more minimal. For while there might be some shred of plausibility in insisting that all evaluative properties correspond to ways of being good *simpliciter*, there is no shred of plausibility in the thought that all standards of correctness correspond to ways of being genuinely normative. Yet for someone engaged in an activity, considerations that bear on whether acting in some way would contribute to satisfaction of the activity's constitutive standard of correctness have the same claim to being reasons as reasons for envy and amusement. Insensitivity to both bears on the appropriateness of certain forms of criticism, where the degree of appropriateness is modulated by how the relevant considerations balance out.

The bottom line, then, is that one should not claim that right-kind considerations and normative reasons are different kinds of things; indeed, as we noted earlier, one cannot make this claim without abandoning the "everybody's problem" response to the WKR problem. A more plausible view is that normative reasons are a type of right-kind reason. But we need

a way to draw the line, other than by gesturing at a special kind of normativity that seems to smuggle in what they are trying to analyse.

3 The Fundamentalist's Dilemma

So far our conclusions threaten everyone who wants to analyze normative notions in terms of normative reasons; given how plausible reasons-based analyses are of certain notions, our conclusions so far threaten just about everyone. Just as everyone has to differentiate between favorers of the right-kind from favorers of the wrong-kind, everyone has to differentiate favorers of the right-kind that are robustly normative from favorers of the wrong-kind that are only anaemically normative. We promised, however, to raise deep worries for fundamentalism. We'll do that now.

We think the RKR problem raises a deep worry for fundamentalists. This is because fundamentalism imposes tight constraints on how the problem can be solved, and these constraints make it quite unclear how fundamentalists can solve the RKR problem at all.

Here is what the fundamentalist must do in order to solve the RKR problem. First, the fundamentalist must tell us which activities and attitudes are normatively relevant. Which ones are such that right-kind reasons for them are robustly normative? Second, the fundamentalist must tell us what grounds the distinction between activities that are normatively relevant and activities that are not—i.e., what these activities have in common in virtue of which they are robustly normative. She should not rest content with a mere list.

These tasks are *particularly* hard for fundamentalists because it seems clear that the fundamentalist cannot appeal to *anything genuinely normative* to solve the problems. After all, there are only two options for a fundamentalist who appeals to normative properties to solve the problem: either she can appeal to something other than normative reasons or she can appeal to normative reasons. To appeal to some non-reason normative property would call fundamentalism into question. It would then be plausible that that property is at least as fundamental as reasons.

To appeal to normative reasons won't help, since what we want, precisely, is a story about which right-kind reasons are genuinely normative reasons. To appeal to normative reasons at this point would be to smuggle genuine normativity in through the back door, seemingly independently of reasons. But part of the job description of right-kind reasons was to analyze such normativity.

At this point things are looking pretty grim for the fundamentalist. It is unclear what resources are available to the fundamentalist to solve the Normativity and Rationale Problems and thus to solve the RKR problem. Without normative stuff, what is left? The story would have to be told in non-normative terms. But it is hard to see what that story could be.

We call this problem the *Fundamentalist's Dilemma*. The only options for solving the RKR problem are to appeal to normative properties or natural properties; the former route seems to be closed off by the constraints of fundamentalism, and it is simply hard to see what the latter

route would involve. Even if the RKR problem doesn't strike one as pressing on its own, these problems lead straightforwardly to the Fundamentalists' Dilemma, which is pressing indeed.

4 Possible Solutions to the Fundamentalists' Dilemma

Is the Fundamentalists' Dilemma crippling? We hope not, since we are ourselves attracted to fundamentalism. In this section, we will consider some possible responses to the dilemma and indicate what we take to be the most promising solution.

Before we consider the options, it will be helpful to situate the options vis-a-vis *reduction*. Unsurprisingly, we have seen that fundamentalists cannot analyze normative RKR's in terms of something else that is normative. To do this would be to give up on fundamentalism. Hence, fundamentalism is incompatible with a certain *type* of reduction. This is a reduction of normative RKR's to something else that is *normative*. Despite this, fundamentalism is compatible with another type of reduction. This is a reduction of normative RKR's to something *descriptive* or *natural*. So the fundamentalist has two options: either they think that normative RKR's can be reduced to something descriptive or natural, or they think that they are irreducible. We will call the latter view *reasons primitivism*.

As you might expect, there is more to say about reductive views than about primitivist views. Primitivist views, in effect, hold that normative RKR's are ground-floor fundamental. They are absolutely basic building blocks. There is not much one can do to positively develop this view. As even its proponents acknowledge, it is to be defended negatively—it is vindicated only by showing that all other views fail.¹⁶ That said, below we will suggest that in fact something positive can be said by the primitivist about the Fundamentalists' Dilemma. Before we get to that, we'll discuss two reductive accounts.

4.1 Humeanism

The Humean maintains that facts about normative RKR's reduce to facts about one's desires. Following Schroeder (2007), we can say that for the Humean, normative reasons for A to ϕ are facts that explain why ϕ -ing promotes at least one of A 's desires. If we combine this with Activity, we can get an analysis of normative RKR's:

Humean Activity: What it is for r to be a normative RKR for A to ϕ is for r to be a reason shared by everyone engaged in a and for r to explain why ϕ -ing promotes some of A 's desires.

Humean Activity is a conjunctive account of normative RKR's. Interestingly, given a *prima facie* appealing assumption about being engaged in activities, the Humean can hold that in fact Activity is the correct analysis of normative RKR's.

¹⁶See Parfit (2011), Scanlon (2014) and especially Korsgaard (1996).

The assumption is this: necessarily, if you are engaged in an activity *a*, then you have at least one desire that will be promoted if you do what everyone engaged in *a* has a reason to do in virtue of being engaged in *a*. If this assumption is true and the Humean view is true, then it turns out that all RKR's are normative reasons. That is, it turns out we were wrong above when we insisted that some standards of correctness are normatively anaemic. For the Humean, any standard of correctness can be normatively full-blooded for some agent *A* just in virtue of *A* engaging in an activity governed by that standard.

By making this move, the Humean fundamentalist doesn't solve the RKR problem, but rather dissolves it. This is because the problem arises only given the claim that not all RKR's are normative reasons. The Humean who makes good on the assumption thinks that all standards of correctness generate normative RKR's for those agents engaged in the relevant activities. So they have no burden explaining which activities generate normative RKR's, nor any burden explaining why the activities that generate normative RKR's do so. They presumably meet the latter burden just in virtue of vindicating Activity.

The ease with which the Humean view dissolves the problems is a serious virtue of the view. This virtue, as far as we know, has never been articulated (although we suspect Schroeder is well aware of it). However, this virtue should carry the day only if we should accept both the Humean view of reasons and the assumption about engaging in activities (which we will call the *Engagement Assumption*). Unfortunately, both are seriously questionable. Furthermore, we think that the triad of views consisting of the Humean view, the Engagement Assumption, and Activity is implausible despite its ability to dissolve the problem. We will defend each of these claims in turn.

Let's start with the Engagement Assumption. While the Engagement Assumption does have initial plausibility, we think that it doesn't stand up to scrutiny. It seems initially plausible because many activities are essentially *intentional*. For those activities, it is plausible that someone engaged in that activity will necessarily have a desire promoted by conforming to the standards of the activity. That desire will be the desire associated with one's intention to engage in the activity.

The problem is that not all activities relevant to RKR's are essentially intentional.¹⁷ Consider believing. One can engage in the activity of believing without intending to so engage. At some point in our development we come to have the capacity to believe, and at that point we just start believing. An intention to have beliefs is not necessary. Thus, it seems that believing is not an essentially intentional activity and so is far from obvious that everyone has a desire that would be promoted by conforming to the constitutive standards of belief in virtue of engaging in it. So the Engagement Assumption doesn't look airtight. In fact, it looks like it undergeneralizes.

Furthermore, it is intuitively plausible that some standards are not genuinely normative.

¹⁷Schroeder concedes this point. He tries to get out of the problem by arguing that for each non-essentially intentional activity, there will be some essentially intentional activity that correlates with it (e.g., with admiration he suggests emulation). Sharadin (2013, 2016) shows that this strategy will not work.

Of course, it is no secret that Humeanism has revisionary implications. And much has been said on both sides of the debate. We don't expect to settle the debate here. Nevertheless, the point can be put in a special way here. Consider certain ritualistic killings. Given that they are ritualistic, there are standards of correctness governing them. It is possible to perform the killings in a way that deviates from these standards. So a WKR problem can arise for these ritualistic killings. Are the RKR's relative to those standards normative RKR's for those (intentionally) engaged in those activities? Plausibly not. This would make engagement in horrid activities bizarrely self-vindicating.

We know that there are possible ways out for the Humean. In fact, we recognize that we have added ammunition to their arsenal by showing that they have a powerful reply to the Fundamentalist's Dilemma. Nonetheless, the triad of views is not all that plausible. This is for two reasons. Firstly, it doesn't look like the Engagement Assumption that is needed to get the reply to work is going to get the right predictions out of Activity. That is, it doesn't look like the kinds of activities relevant to Activity are essentially intentional activities. But that is what is needed in order for the Humean view to show that all RKR's are normative reasons, which is what the Humean needs in order to dissolve the Fundamentalist's Dilemma. The second problem is a deeper problem with Humeanism: it generates too many reasons and generates them too easily.

4.2 Constitutivism

A second reductive response is a version of constitutivism. It attempts to solve the Normativity and Rationale Problems in one fell swoop by telling us that the activities for which right-kind reasons are normative reasons are the activities in which any agent must be engaged in *qua* agent. It defuses the Fundamentalist's Dilemma by noting that the line it draws to distinguish between right-kind reasons that are normative and right-kind reasons that are not is a descriptive line. In this way, constitutivism provides an elegant reductive solution to the fundamentalist's problems.

Let's walk through the solution at a slower pace. Let's start with the Normativity Problem. The constitutivist might start by pointing out that not all activities are created equal. In particular, some activities seem essential to our *agency*. We can call these activities the activities essential to *deliberation*. When we deliberate practically, we deliberate about what to want, what to intend, and what to do. When we deliberate theoretically, we deliberate about what to believe and how confident we should be about various claims. To engage in deliberation is *just what it is* to be an agent. This is a natural or descriptive claim about what agency is.

The constitutivist holds that we can reduce facts about normative RKR's to facts about agency. In particular, we can analyze normative RKR's in terms of the activities that are constitutive of agency:

Constitutive Activity: What it is for r to be a normative RKR to ϕ relative to some activity a is for r to be a reason shared by everyone engaged in a in virtue of being engaged in a and for a to be an activity constitutive of agency.

This is a constitutivist view twice over because it understands RKR in terms of the constitutive norms of activities and it understands *normative* RKR in terms of what's constitutive of agency. We will call the former view *constitutivism about RKR* and the latter position *constitutivism about normative RKR*.¹⁸

This is an interesting analysis. It is also predictive once it is supplemented with views about the activities constitutive of agency. And there are already prominent views in the literature about particular activities, with believing being the most well-developed.¹⁹ The plausibility of these views provides indirect evidence for Constitutive Activity. Consider belief again. Believing is plausibly on the list of normatively relevant activities. It is plausible that right-kind reasons for belief are normative reasons. The standard of correctness for belief is, one might think, surely not on all fours with norms of etiquette or some specific set of instructions for cooking *cacio e pepe*. But what could distinguish the standard of correctness for belief from these other standards other than the fact that “playing the belief game” is a constitutive feature of agency? If the constitutivist story provides a credible answer—perhaps the only obvious answer—to this question about the normativity of belief's standard of correctness, explanatory parsimony recommends telling the same story elsewhere.

This version of constitutivism has an obvious solution to offer to the Rationale Problem. What unifies the activities on the list is the fact that they are all constitutive of agency. One might find appeals to agency normatively irrelevant, but all should agree that this type of appeal to what's constitutive of agency is interestingly unified.

There are, of course, challenges for constitutivists. Our goal isn't to provide a comprehensive defense of constitutivism here, but we will consider two challenges. The first is the most notable general challenge to constitutivism and is often thought to plague all versions of the view. This is the “schmagency” challenge.²⁰ This challenge, though, does not obviously apply to the type of constitutivist view appealed to here. This is because the force of the challenge is felt only when certain background assumptions are made about *authority*. These background assumptions are made by some constitutivists. The most notable example is Christine Korsgaard.²¹ Korsgaard is interested in employing constitutivism in a particular context. The operative question in this context is a question about the normative authority of morality—why be moral? Given Korsgaard's unorthodox views about the function of moral concepts, she has a very demanding view about what it takes to vindicate the normative authority of morality.²² Any satisfactory answer, for Korsgaard, will strike any agent inquiring about what to do as a satisfactory answer. This standard demands from a successful answer that it convinces a skeptic—a tall order indeed.

¹⁸We were implicitly relying on constitutivism about RKR when we laid out Activity above. We are making it explicit here mostly because it will be helpful in the next subsection.

¹⁹See, e.g., Velleman (2000a), Shah (2003), Sosa (2015), Wedgwood (2017). For discussion of intention, see Shah (2008).

²⁰Originally articulated in Enoch (2006). See also Enoch (2011).

²¹See especially Korsgaard (1996).

²²See (Korsgaard, 1996, ch. 1) for the (complicated) details.

These background views open Korsgaard up to a particularly damning version of the schmagency problem. For it seems entirely intelligible to skeptically challenge the normative relevance of what's constitutive of agency. The normative relevance of agency is an open question like (nearly) anything else. One can intelligibly ask why one shouldn't opt out of agency in order to become a schmagent. What this shows is that Korsgaard's demands for a successful answer cannot be met. She has a schmagency problem.

But not all constitutivist views have the ambitions of Korsgaard's view. We are assuming that the proponent of Constitutive Activity is not in the business of convincing the skeptic *via* Constitutive Activity. Moreover, we recommend to the proponent of Constitutive Activity that she not go in for answering whether there is some special reason for being an agent rather than a schmagent. That question has a false presupposition—viz. that constitutivism is committed to thinking that there is such a reason.

Sometimes we have no reasons to be agents any longer. Sometimes we should kill ourselves. Other times we might not have particularly compelling reasons to be an agent, but, nevertheless, we are agents. We thus have all sorts of normative reasons, *per* Constitutive Activity. The fact that, in these cases, agents don't have particular strong reasons to be agents is neither here nor there. Without appealing to claims about authority, it is hard to generate pressure to think otherwise. The constitutivist should deny those overly ambitious claims about authority.²³

This challenge also strikes us as weak in the context of the Fundamentalists' Dilemma. Reasons fundamentalism had a lot going for it. It offers a beautiful account of the internal structure of the normative. Constitutivism provides the obvious resolution of the Fundamentalists' Dilemma. It is plausible on first-order grounds that the list of activities suggested by constitutivism is the list of activities for which right-kind reasons are normative reasons. But if one is already attracted to reasons fundamentalism, there seems to be no deep explanation of why that list is the right list other than the explanation given by the constitutivist. If reasons fundamentalism is true, the Fundamentalists' Dilemma must have an answer. Perhaps the best argument for constitutivism is that it provides this answer.

The second challenge we will consider is a more local worry for Constitutive Activity. The core of the worry is that Constitutive Activity undergenerates because not enough activities are constitutive of agency to generate enough normative RKR's. We can see the worry by thinking about particular attitudes. Admiration seems to be the sort of attitude we have normative RKR's to have. According to Constitutive Activity, this is true only if the activity that governs admiration is an activity that is constitutive of agency. But this is far from clear. It seems plausible that we would be the same sort of creature even if we didn't have the capacity for admiration.

We think that this is a serious challenge to Constitutive Activity. Further, the point is obviously not limited to admiration. There are likely many other reactions that intuitively are reactions we have normative RKR's to have even though it is not obvious that the activi-

²³For a similar response, see [Smith \(2015\)](#).

ties associated with those reactions are constitutive to agency. We'll call these Normative But Optional activities. We can't hope to meet these challenges here, but we will sketch out two strategies for meeting them.

The first strategy is simple. It insists that some particular Normative But Optional activity is in fact constitutive of the sort of agents we are. We will call this the Doubling Down strategy. In order to apply the Doubling Down strategy in all cases, it seems to us that one would need to adopt a fragile conception of what constitutes the sort of agency normal agents exemplify. On such a fragile understanding, it is easy to cease being the sorts of agents normal agents are. In fact, it might be that some actual human beings fail to be the kind of agents that normal agents are (e.g., perhaps psychopaths fail to be the sort of agents normal adults are). Such a fragile view will have some implications that some will find unappealing, but they don't strike as wildly implausible on their face.²⁴

Furthermore, one needn't accept the fragile view of agency in order to use the Doubling Down strategy in a more local way. For example, we are actually skeptical that admiration is a Normative But Optional activity even if one accepts a robust account of the constitution of our agency. This is because we find it plausible that admiration is a central reaction to our moral agency, which is a constitutive part of our agency even if you have a robust view.²⁵

The second strategy insists that the purported Normative But Optional activities are reducible to reactions that are obviously constitutive of our agency, with the obvious contenders being belief and desire. We call this the Reduction Strategy. In order to apply the Reduction Strategy in full generality, one would likely commit oneself to a sparse philosophical psychology. We note, though, that the exact sort of project that seems needed is already rigorously defended by several prominent philosophical psychologists—e.g., [Arpaly & Schroeder \(2014\)](#), [Sinhbabu \(2017\)](#).

Of course, just like it was with the Doubling Down strategy, one might apply the Reductive Strategy more locally. When it comes to admiration, for example, one might argue that admiration reduces to a bundle of beliefs and desires. To take an oversimplistic view, one might hold that admiration is the bundle of a belief that the object of admiration is desirable and the desire to emulate the admired (for discussion, see [Zagzebski \(2015\)](#)). If such a reduction were true (or a suitably more complex version), then it would be plausible that RKR for admiration are normative RKR. For, if the reduction were true, then RKR for admiration would just be RKR for having certain beliefs and desires. Believing and desiring are clearly constitutive of agency and thus RKR for believing and desiring are normative, per Constitutive Activity.²⁶

²⁴Some evidence for this is that Korsgaard adopts a very fragile view of one's practical identity in recent work (see especially [Korsgaard \(2009\)](#)).

²⁵See [Zagzebski \(2015\)](#) and [Irwin \(2015\)](#).

²⁶Interestingly, this doesn't quite show that *admiration* is constitutive of agency. For even if the reduction is true, it seems possible to be the type of agents we are without having admiration—i.e., without having the bundles of beliefs and desires that constitute admiration. Even if it doesn't show that admiration is constitutive of agency, it does show that the RKR for admiration are normative RKR for those that do admire. This might

These are the strategies we think defenders of Constitutive Activity should use to combat cases of purported Normative But Optional activities. While we are skeptical that either can be used in full generality, we are hopeful that local uses of each strategy can stave off the most pressing worries. What we've hoped to show here is that with more work these objections can be overcome and thus that we should continue to think of Constitutive Activity as a live option.

While this argument on its own might not sell constitutivism, it in combination with reasons fundamentalism and the need for an answer to the Fundamentalists' Dilemma strike us as providing a serious argument for constitutivism. At the very least, the disjunction *either constitutivism is true or reasons fundamentalism is false* merits high credence.

4.3 A Louder Quietism

We will end by discussing a non-reductive view. We end with this view not because we think it is the most plausible—far from it—but rather because our discussion of the reductive views will help us see how a quietist position can be a bit more informative than some quietist views.²⁷ Just to be clear, however: this quietist view is a quietist view. That is, it holds that there is nothing informative that can be said at the analytic level about why the normative RKR's are normative RKR's. They cannot be analyzed in terms of something normative—hence fundamentalism. They cannot be analyzed in terms of something natural or descriptive—hence primitivism. This much is just run-of-the-mill quietism.

Nevertheless, according to the louder quietism we have in mind, the RKR problem does have a solution of a certain kind. There is a compact list to be given of the attitudes for which right-kind reasons are normative reasons, and there is a principled story to be told about what the members of this list have in common. But according to our quietist, the principled story to be given is at bottom a piece of *first-order normative theorizing*, one that is supported by systematizing our intuitions about what is a reason for what. It doesn't tell us anything about how the normative is grounded in the natural or descriptive, *pace* the constitutivist and Humean.

To make this view more concrete, we will sketch one way that it could be developed, though it is only one that we mention for the sake of illustration. This version of the view gives the same list of normatively relevant activities that the constitutivist view gives in solving the Normativity Problem, and even agrees with the constitutivist about what the elements of that list have in common. But it provides a first-order rationale for the latter claim. The most

seem to threaten the letter of Constitutive Agency. We don't think that it does. The right way to think of it is that if the reduction is true, then the activities that are relevant are believing and desiring. So if admiration isn't constitutive but the reduction is true, admiring is just not the activity that we plug into the analysis in order to determine what the normative RKR's are for admiring. We need to plug in believing and desiring. Thanks to an anonymous referee for prompting us to say more here.

²⁷We have in mind Parfit and Scanlon in particular. We recognize that some quietists might be immune to everything we say below. We also recognize that Parfit and Scanlon have engaged in the sort of first order project we discuss when it comes to parts of the normative domain, with morality being the domain they have focused the most on. The view we sketch here is supposed to play a similar role but at much more general level.

natural way for quietists to do all of these things is by becoming constitutivists about RKR. That is, by thinking that the reasons that are shared by everyone engaged in an activity in virtue of being engaged in that activity are the reasons that have some connection to the constitutive standards of that activity. Normative reasons are then taken to be RKR generated for the activities constitutive of agency, where this is made as a first-order claim. It is a first-order view with some plausibility, as the literature on how to derive a truth-norm on belief by appeal to constitutive norms illustrates.²⁸

To see the main point more clearly, think about the Moorean view about goodness and pleasure. The Moorean view about goodness is a primitivist and non-naturalist one. It holds that goodness is an unanalyzable property that is different in kind from the natural properties. Nevertheless, the Moorean thinks that the only states of the world that had this property are pleasure states. This doesn't mean that the property of being good is the same as the property of being pleasant. Rather, it's just to say that *the good* is pleasure. So a certain kind of first-order hedonism is, of course, compatible with primitivist non-naturalism.

The non-naturalist constitutivist view we've imagined is like the Moorean version of hedonism. It holds that the right-kind reasons (the reason analogue of the good) are the facts that are tied to constitutive standards in a certain way. This doesn't mean that the property of being a normative reason is the same as the property of being a fact that stands in the right relation to constitutive standards. All that is posited is a necessary correlation.²⁹

This version of quietism is more interesting than run-of-the-mill quietism because it gives a unifying view about which facts are normative RKR. It doesn't give us a metaphysical explanation of *why* they are like this. Nor does it tell us that the *kind of things* RKR are is the same as the kind of things tables, chairs, and electrons are. Nevertheless, it tells a unifying story about which facts are the normative RKR.

While this view has a bit more predictive power than run-of-the-mill quietism, it is still anchored in a strange bit of metaphysics. It posits a *brute* necessary correlation between RKR and constitutive standards.³⁰ Furthermore, this an essential feature of the view. This is what

²⁸See, e.g., Shah (2003), Velleman (2000b).

²⁹The existence of Moorean hedonism doesn't show that there aren't intelligible metaphysical questions to ask about what the good *is*. Of course, Mooreans deny that these questions have substantive answers. But it would be a bad move to claim that there are no intelligible metaphysical questions to ask about the nature of goodness simply because one can hold that the independent property of goodness necessarily correlates with pleasure. Analogously, it would be a bad move to hold that we cannot intelligibly ask the metaphysical question of what RKR are normative RKR just because one could hold that the independent property of being a normative RKR necessarily correlates with the property of being a consideration that is a reason shared by all engaged in an activity constitutive of agency.

³⁰One might wonder if this only looks odd because we have been assuming that the solution to the Normativity Problem maintains that all the RKR for the robustly normative activities are normative RKR. If one were to deny this, would there be a brute necessary connection? Well, it is right that the brute connection we are complaining about here is a brute connection between normative RKR and the considerations that provide evidence that the standards of correctness of certain activities constitutive of agency are met. But even quietists who deny that all the RKR of constitutive activities provide normative RKR posit a brute connection. Any quietist who holds that there is a correlation between some descriptive facts and some normative

makes it a quietist view. This makes it stark that the ‘solution’ it offers to the problems isn’t much of a solution. The best the quietist can do is change the subject from the second-order question about the metaphysics of normative RKR to the first-order question about what reasons there are. They remain silent about the second-order question except to say that it is not a question worth trying to answer.

4.4 Sizing Up the Options

We’ve explored three replies to the Fundamentalists’ Dilemma. The first two were reductive. They attempt to analyze normative RKR in terms of something natural. This is compatible with fundamentalism because fundamentalism is merely the claim that normative RKR are *normatively* fundamental. The last view was a quietist non-naturalist view. It held that normative RKR are not analyzed at all. They are basic building blocks. Nevertheless, we saw that the quietist could borrow some machinery from the reductive constitutivist account in order to be have a bit more predictive power.

We think the best option is the reductive constitutivist one. Constitutivism about RKR is, we think, the most natural way of developing Activity. Moreover, it is striking how plausible it is that the activities that are constitutive of agency are the activities that would naturally go on one’s list of activities that are normative in a full-blooded sense. This is obviously far from decisive, but we are intrigued by it.

Furthermore, the constitutivist view has serious virtues compared to the other two views. It has a principled way of avoiding the explosion of normative RKR that the Humean is saddled with. The constitutivist view lets us hold onto the plausible thought that not all standards are created normatively equal. Some of them are normatively irrelevant for everyone. The Humean, by contrast, holds that any standard can become normatively relevant for an agent just in virtue of that agent engaging in an activity governed by those standards. That makes full-blooded normativity too cheap.

The constitutivist view has a serious explanatory advantage over the quietist view. The quietist view posits a brute metaphysical necessity between normative RKR and the constitutive standards of the activities that are constitutive of agency. Brute metaphysical necessities are born of desperation. We should avoid desperation if we can. The reductive constitutivist view allows us to avoid desperation because it offers an explanation of the metaphysical necessity. There is a necessary correlation between normative RKR and the constitutive standards of the activities constitutive of agency because *what it is* to be a normative RKR is to be a RKR relative to an activity that is constitutive of agency.

These are serious virtues of the constitutivist view. They are obviously not fully decisive. We haven’t surveyed all the possible solutions to the Fundamentalists’ Dilemma. Nor have we

facts posits a brute connection. So you can’t get out of a brute connection by giving up our assumption. By giving it up, you will certainly have a less elegant theory, though. Pending a surprising explanation of why some RKR associated with constitutive activities are not normative, this view seems to lose the elegance of the view we consider in the text. Thanks to an anonymous referee for raising this point.

seriously considered the option of abandoning fundamentalism because of the dilemma. Our interest has been in clearly articulating the dilemma and seeing if there are good routes out of the problem. The constitutivist route is the best option we've thought of. Whether this a failure of our imaginations is yet to be determined.

5 Concluding Morals

Let's take stock of the key morals. The first key moral is that the generality of the WKR/RKR distinction shows that there is another problem that is structurally similar to the WKR Problem. This is the RKR Problem. The generality of the WKR/RKR distinction gives rise to the RKR problem because the generality of the problem shows that not all RKR's are *normative* reasons. Accordingly, those of us who analyse complex normative notions in terms of normative reasons—i.e., everyone—cannot pick out our analysans just by providing a principled distinction between WKRs and RKR's. We also need a story about which activities/attitudes are the ones for which RKR's are normative reasons. But this task is highly non-trivial. Further, this task is especially difficult for the fundamentalist. We demonstrated this via the Fundamentalist's Dilemma. Fundamentalism blocks one from appealing to anything normative to solve the RKR problem. This forces the fundamentalist to pick between quietism and reductionism.

The moral of the story is that to properly solve the RKR problem, the fundamentalist is forced into taking a stand about the metaphysical nature of normative reasons. That is, she has to take a stand about whether naturalism or non-naturalism is true. In some ways this is unsurprising once you go in for the thought that *something* needs to be said about what the basic building blocks are like. And the WKR problem clearly forces the fundamentalist to say something about this. Relying on the 'counts in favor' idiom doesn't cut it. But it's obvious that the fundamentalist cannot appeal to anything normative to explain which facts that count in favor are her building blocks. So if she is going to say anything with explanatory power, she will need to appeal to something natural.

The surprising twist is that she will have to do more than solve the WKR problem. She'll also have to solve the RKR problem. As we saw, this problem breaks down into the problem of telling us which activities/attitudes are the ones for which RKR's are normative reasons, and the problem of telling us why. It is here, we suggested, where the fundamentalist will take a stand over naturalism. We have suggested that the stand she should take is in favor of a form of reductive constitutivism. This view offers the best balance of explanatory power and extensional adequacy. And those are the right kind of reasons for accepting a theory.

References

Arpaly, N. & Schroeder, T. (2014). *In Praise of Desire*. Oxford University Press.

Baker, D. C. (FC). Skepticism about ought simpliciter. *Oxford Studies in Metaethics*, 13.

- Copp, D. (2004). Moral naturalism and three grades of normativity. In *Normativity and Naturalism*. Ontos-Verlag.
- Dancy, J. (2004). *Ethics without Principles*. Oxford University Press.
- D'Arms, J. & Jacobson, D. (2000a). The moralistic fallacy: On the "appropriateness" of emotions. *Philosophy and Phenomenological Research*, 61(1), 65–90.
- D'Arms, J. & Jacobson, D. (2000b). Sentiment and value. *Ethics*, 110(4), 722–748.
- Enoch, D. (2006). Agency, shmagency: Why normativity won't come from what is constitutive of agency. *The Philosophical Review*, 115(2), 169–198.
- Enoch, D. (2011). Shmagency revisited. In M. Brady (Ed.), *New Waves in Metaethics*. Palgrave-Macmillan.
- Foot, P. (1972). Morality as a System of Hypothetical Imperatives. *The Philosophical Review*, 81(3), 305–316.
- Irwin, T. H. (2015). Nil Admirari? uses and abuses of admiration. *Aristotelian Society Supplementary Volume*, 89(1), 223–248.
- Korsgaard, C. (1996). *The Sources of Normativity*. Harvard University Press.
- Korsgaard, C. M. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press.
- McPherson, T. (2011). Against quietist normative realism. *Philosophical Studies*, 154(2), 223–240.
- Parfit, D. (2011). *On What Matters*. Oxford University Press.
- Plunkett, D. & Shapiro, S. (2017). Law, morality, and everything else: General jurisprudence as a branch of metanormative inquiry. *Ethics*, 128(1), 37–68.
- Rabinowicz, W. & Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting proattitudes and value. *Ethics*, 114(3), 391–423.
- Scanlon, T. (2014). *Being Realistic about Reasons*. Oxford University Press.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Schroeder, M. (2010). Value and the Right Kind of Reason. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, volume 5. Oxford University Press.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, 112(4), 447–482.

- Shah, N. (2008). How action governs intention. *Philosophers' Imprint*, 8(5), 1–19.
- Sharadin, N. (2013). Schroeder on the wrong kind of reasons problem for attitudes. *Journal of Ethics and Social Philosophy*, 7, 1–8.
- Sharadin, N. (2016). Reasons wrong and right. *Pacific Philosophical Quarterly*, 97(3), 371–399.
- Sinhababu, N. (2017). *Humean Nature: How Desire Explains Action, Thought, and Feeling*. Oxford University Press.
- Skorupski, J. (2011). *The Domain of Reasons*. Oxford University Press.
- Smith, M. (2015). The magic of constitutivism. *American Philosophical Quarterly*.
- Sosa, E. (2015). *Judgment and Agency*. Oxford University Press.
- Velleman, D. (2000a). On the aim of belief. In D. Velleman (Ed.), *The Possibility of Practical Reason*. Oxford University Press.
- Velleman, D. (2000b). *The Possibility of Practical Reason*. Oxford University Press.
- Wedgwood, R. (2017). *The Value of Rationality*. Oxford University Press.
- Wodak, D. (FC). Fictional normativity and normative authority. *Canadian Journal of Philosophy*. Manuscript, Virginia Tech University.
- Woods, J. (FC). The authority of formality. *Oxford Studies in Metaethics*, 13.
- Zagzebski, L. (2015). I—admiration and the admirable. *Aristotelian Society Supplementary Volume*, 89(1), 205–221.