

primary reasons as normative reasons

Pre-review version

Please reference final version forthcoming in Journal of Philosophy

1. Introduction

Donald Davidson's conception of motivating or "primary" reasons for action as belief-desire pairs was and remains hugely influential in action theory.¹ But it does not enjoy comparable influence in moral theory.² Primary reasons are often treated as categorically distinct from moral reasons, and from normative reasons more generally, usually on the grounds that moral reasons (i) can be motivating reasons and (ii) are facts. This reasoning underlies an orthodoxy according to which both motivating and normative reasons are facts, propositions, or states of affairs.³

But the thesis that normative reasons for action are primary reasons enjoys better support than its critics imagine. This paper defends this thesis, proceeding as follows. I begin by showing that a popular view about what makes reasons normative implies that normative reasons are primary reasons because primary reasons draw crucial distinctions in our motives that facts generally don't. I'll then show that because primary reasons draw crucial distinctions between kinds of justification that facts generally don't, normative reasons are primary reasons. Finally, I'll describe an important but underappreciated problem for the orthodox conception of normative reasons and show that it does not arise for the competing Davidsonian conception. These three claims form the basis for thinking of normative reasons for action as primary reasons.

The primary reasons I discuss differ in one important respect from the ones that Davidson describes, at least on one reading. According to that reading, primary reasons are pairs of tokened beliefs and desires. For example, when I intentionally drink water, my primary reason, on this view, is my belief that water will quench my thirst paired with my desire to quench my thirst. This conception of primary reasons prevents them from being normative reasons. That's because agents can have a normative reason for acting without tokening any of the belief-desire pairs relevant to that action. This happens when agents are not moved to do what they morally should. So normative reasons cannot be pairs of beliefs and desires.

¹ See his "Actions, Reasons, and Causes," *The Journal of Philosophy* 60.23 (1963): pp. 685-700. I'm using 'desire' in Davidson's broad sense to denote a wide range of pro-attitudes.

² There are exceptions. For example, Robert Myers defends the thesis that desires relate us to normative reasons in "Desires and Normative Truths: A Holist's Response to the Sceptics," *Mind* 121 (2012), pp. 375-406. Paul Hurley defends a related view in "Desire, Judgment, and Reason: Exploring the Path Not Taken," *The Journal of Ethics*, 11.4 (2007): pp. 437-463.

³ Prominent recent advocates of this view include, *inter magna alia*, Jonathan Dancy in *Practical Reality* (Oxford University Press, 2000), Mark Schroeder in *Slaves of the Passions* (Oxford University Press, 2007), Daniel Star in *Knowing Better* (Oxford University Press, 2015), and Errol Lord in *The Importance of Being Rational* (Oxford University Press, 2018).

2. Primary Reasons and Good Reasoning

Rather, I'll instead assume that primary reasons are the contents of those belief-desire pairs. On this account, when I intentionally drink water, my primary reason is the content of the belief that water will quench my thirst paired with the content of desire to quench my thirst, *i.e.*, the *fact* that water will quench my thirst and the *goal* of quenching my thirst. So, on the view I assume, the primary reasons eligible to be normative reasons are fact-goal pairs. What gives this view continuity with Davidson's is that its primary reasons are the contents of Davidson's reasons.

While this paper focuses on advantages for moral theorizing that come from thinking of normative reasons as fact-goal pairs, this conception of reasons may also circumvent a well-known challenge to Davidson's programme in the philosophy of action. It is often assumed that reasons are what we deliberate about when we act intentionally.⁴ However, as several philosophers have observed,⁵ our psychological states are not normally the things *about* which we deliberate; they are merely the things *in virtue of* which we deliberate. So motivating reasons, these philosophers conclude, cannot be primary reasons if primary reasons are psychological states. However, thinking of primary reasons as the contents of belief-desire pairs, rather than as the beliefs and desires themselves, dispenses with this worry about psychologism while remaining recognizably Davidsonian. As a result, there are independent grounds for thinking of motivating reasons as the contents of belief-desire pairs. However, I won't dwell on this discussion, turning instead to the titular claim that normative reasons should be thought of as primary reasons.

Two questions about normative reasons easy to elide. The first concerns what *makes* or *grounds* something's status as a normative reason. The second concerns what *kinds of things* can enjoy that status, that is, what things are eligible to be normative reasons. To illustrate, note that the question of what makes someone a MLB baseball player differs from the question of what MLB baseball players are. What makes someone a MLB baseball player is signing a MLB contract. That's an answer to the first question but applied to baseball players, not reasons. However, only *people* are baseball players. That's an answer to the second question. Moral philosophers tend to focus on the first question about reasons. But this paper is principally concerned with the upshots of a Davidsonian answer to the second question. As a result, the call to think of primary reasons as normative reasons is a call to consider a new *model* for theories of normative reasons, for the model of reasons we assume shapes our downstream inquiry about which reasons are genuinely normative and about what makes them normative.

Some accounts about what makes a reason normative, the first question, are neutral on the second question of what things are reasons. For example, Derek Parfit and T. M. Scanlon's position that the normativity of reasons is primitive -- that *nothing* grounds something's status as a reason --

⁴ Schroeder (2007) calls this 'the deliberative constraint' in *op. cit.*, p.26.

⁵ Thomas Nagel offers this criticism in *The View from Nowhere* (Oxford University Press, 1986), pp.141-3. Jonathan Dancy continues this line of criticism in *Practical Reality*, chapter four.

implies no constraints on what reasons are, all by itself.⁶ That's not to say that Parfit or Scanlon are silent on the second question of what reasons are. On the contrary, they appear to agree that reasons are facts. Rather, the point I wish to highlight is that the grounds for their agreement is independent of their account of what makes reasons normative. As a result, their primitivism is compatible with the thesis that normative reasons are primary reasons.

By contrast, some competing accounts of what makes reasons normative *do* impose constraints on what things are reasons; that is, their distinctive answer to the first question constrains their answer to the second. These accounts are of particular interest here for some such constraints make it especially appealing to think of normative reasons for action as primary reasons. For example, Bernard Williams famously argued that there is a normative reason for you to do something only if that reason is, in some sense, be capable of motivating you to act in that way through deliberation. Some accounts hold that this deliberative role of normative reasons is the key to answering the first question about reasons, arguing, broadly, that a reason is normative just when and because it figures in an episode of good reasoning.⁷ In particular, some think that a reason for action is normative just when and because it is a *good basis* for acting that way.⁸ Call this *the good basis* account.

Unlike the primitivist account, the good basis account straightforwardly constrains what kinds of things can be reasons. Clearly, only bases for action can be *good* bases for action -- and our bases for action are our motivating reasons. Consequently, the good basis account implies that all normative reasons are motivating reasons and, as a result, if the account is true, then conclusions about the ontology of motivating reasons settle questions about the ontology of normative reasons.

This connection between kinds of reasons provides the first argument for thinking that normative reasons are primary reasons. The argument begins with the observation that since primary reasons are individuated not *just* by the facts that they contain but also by their goals, primary reasons are strictly more finely grained than the corresponding facts. This fineness of grain makes primary reasons appealing motivating reasons because, as I'll now show, claims about our motives are similarly finely-grained. For example, it seems plausible that the Cubs fan and the Cubs hater can head to Wrigley only because the Cubs are playing. Despite responding to the same fact, the two baseball fans' goals in going to Wrigley differ: one goes to Wrigley to see the Cubs win; the other

⁶ See T. M. Scanlon's *What We Owe to Each Other* (Harvard University Press, 1998) and Derek Parfit's *On What Matters Volumes 1 and 2* (Oxford University Press, 2011).

⁷ For example, in *Reasons Without Rationalism* (Oxford University Press, 2010), Kieran Setiya argues, roughly, that a consideration gives a normative reason for an agent to act in some way just when, given the agent's psychology, being moved to act in that way by that consideration is a 'good disposition of practical thought' (77). Matthew Silverstein argues in "Reducing Reasons" in *Journal of Ethics and Social Philosophy* 10.1 (2016): 1–22 that this condition is met just when the considerations figure in 'sound reasoning'. Jonathan Way argues in "Reasons as Premises of Good Reasoning," *Pacific Philosophical Quarterly* 98.2 (2017): 2051-270 that reasons to ϕ are normative when they're 'appropriate premises for reasoning towards ϕ -ing'. Hille Paakkunainen defends the view that normative reasons are premises in good practical reasoning in her "Can There Be Government House Reasons for Action?" in *Journal of Ethics and Social Philosophy* 12.1 (2017): 56–93.

⁸ Alex Gregory defends this view in "Normative Reasons as Good Bases," *Philosophical Studies* 173.9 (2016): 2291-2310.

goes to see them lose. However, if motivating reasons are simply facts, then we're forced to concede, counterintuitively, that the hater and the fan go to Wrigley for the same reason because they act on the basis of the same fact. But that seems clearly wrong -- after all, one goes to Wrigley to see the Cubs win and the other goes to see the Cubs lose.

This argument suggests that motivating reasons are more finely-grained than facts for, as in the Cubs case, a single fact can provide two distinct motivating reasons. But we can appropriate one of Davidson's observations to support the stronger conclusion that motivating reasons are more finely-grained than facts in the manner of primary reasons. Davidson writes,

A primary reason consists of a belief and a pro-attitude, but it is generally otiose to mention both. If you tell me you are easing the jib because you think that will stop the main from backing, I don't need to be told that you want to stop the main from backing; and if you say you are biting your thumb at me because you want to insult me, there is no point in adding that you think that by biting your thumb at me you will insult me.⁹

In the example, when you ease the jib, we can say that the reason you're easing the jib is *that* it will stop the main from backing *or* we can say that the reason you're doing so is *to* stop the main from backing. On a Davidsonian picture, these two claims attribute the same motivating reason. One uses a *that*-clause and one uses a *to*-clause, each of which refers to one of the reason's two parts. The first refers to the part corresponding to the belief that easing the jib will stop the main from backing. The second refers to the part corresponding to the desire to stop the main from backing. Nevertheless, it's natural to run these claims together for we can pick out this single reason by naming either of its parts. Indeed, it's often excessive to use both claims, as Davidson observes. That these different claims pick out different parts of the same reason explains why we often elide them unless we're careful.

The Cubs fan and the Cubs hater go to Wrigley for different reasons. We can identify these distinct reasons given by the fact that the Cubs are playing at Wrigley through *to*-clauses: the Cubs fan's reason for going to Wrigley is *to see the Cubs win*; the Cubs hater's reason for going to Wrigley is *to see the Cubs lose*. The cogency of these two claims and the fact that they appear to name goals, the contents of the desire to see the Cubs win and the desire to see the Cubs lose, respectively, offers independent reason for doubting that motivating reasons are just facts. Although the orthodox view predicts that if two agents act on the basis of the same fact, then they act for the same reason, thinking of motivating reasons as primary reasons gives motivating reasons their intuitive fineness of grain. This is good grounds for thinking that all possible motivating reasons are primary reasons. And if all normative reasons are possible motivating reasons, as the good basis account implies, then we should think that all normative reasons are primary reasons as well.

⁹ Davidson (1963), *op. cit.*, p.688.

Advocates of the orthodoxy can resist this argument in at least three ways. First, they may insist that goals are facts and that the case is, therefore, consistent with the orthodox view. In other words, they may insist that agents who act for different goals on the basis of different facts and thereby act for different reasons. But this response fails because goals are not facts. The most decisive argument for the distinction between goals and facts involves wading into the muddy waters of the semantics of *to*-clauses: if *to*-clauses do not denote propositions, as some theorists think,¹⁰ then they cannot express facts. Though I agree with these theorists that *to*-clauses do not denote propositions, I haven't the space to explore that rebuttal here.

A simpler argument for the distinction between facts and goals is that goals don't need to be true to give normative reasons, but propositions need to be true, that is, to be *facts*, to give normative reasons.¹¹ For example, if I desire *to* drink water, some think that said desire is equivalent to the desire *that* I drink water, assessed in some *de se* or first-personal mode. However, I can desire to drink water when I'm not drinking water. In that situation, assuming that the contents of desires are propositions, the content of my desire to drink water is not true. Nevertheless, the goal of drinking water gives me a prudential reason to go to the water fountain when I'm thirsty. Consequently, goals needn't be truths to give reasons, so goals don't play the role of facts in normative reasons. A final rebuttal rests in observing that facts and goals move us differently.¹² For example, it is widely assumed that facts cannot move us as motivating reasons unless we believe them.¹³ But, even if goals were facts, we needn't bear a *cognitive* relation to them when they motivate us; only a *conative* one. Consequently, there's an important motivational difference between facts and goals for which with the orthodox view does not account. This is reason enough to look beyond the view.

¹⁰ I am sceptical that non-finite clauses like 'to drink water' denote propositions when embedded in reasons attributions, largely for reasons laid out in Gennaro Chierchia's "Anaphora and attitudes *de se*" in *Semantics and Contextual Expressions*, edited by Bartsch, van Benthem & van Emde Boas (Kluwer & Reidel, 1989), which trace back to David Lewis's "Attitudes *De Dicto* and *De Se*," *The Philosophical Review*, 88.4 (1979): 513-543. Theorists have since developed this position, providing both syntactically- and semantically-grounded resistance to the claim that *to*-clauses denote propositions. For example, Lucy Campbell's "Propositionalism about intention," *Canadian Journal of Philosophy* 49 (2019): 230-252, especially pp. 241-7, and Nate Charlow's "Metasemantic Quandaries" in *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*, edited by Dunaway and Plunkett (University of Michigan, forthcoming), each argue that syntactic arguments for thinking that *to*-clauses express covert pronouns are undermotivated. Likewise, David Plunkett, Howard Nye, and John Ku's "Non-Consequentialism Demystified," *Philosophers' Imprint* 15.4 (2015): 1-28 and Michael Milona and Mark Schroeder's "Desiring Under the Proper Guise" in *Oxford Studies in Metaethics Vol. 14*, edited by Russ Shafer-Landau (Oxford, forthcoming), identify semantic differences between what *to*-clauses and closely related sentential clauses express.

¹¹ Some, like Schroeder (2007), *op. cit.*, distinguish between subjective and objective normative reasons, allowing that the former can be false propositions. Aficionados of this view should read my claims as implicitly restricted to objective normative reasons.

¹² Pettit and Smith contrast the different deliberative roles played by beliefs and desires in "Backgrounding Desire," *Philosophical Review* 99.4 (1990): 565-592, but the parallel point holds for facts and goals. Mark Johnston's "The Authority of Affect" in *Philosophy and Phenomenological Research* 63.1 (2001): 181-214 also illuminates desire's background role in deliberation.

¹³ Chapter four of Lord (2018), *op. cit.*, offers a comprehensive discussion of this claim.

However, philosophers can resist this argument two other ways, by denying either of two other claims on which it rests. They may deny that all normative reasons are motivating reasons.¹⁴ Or they may deny that the Cubs fan and the Cubs hater go to Wrigley for different reasons, finding it natural to think that they go to Wrigley for the same reason precisely because they go on the basis of the same fact. Fortunately, these two responses can be rebutted by a distinct but closely related argument that does not depend on either claim.

3. Distinguishing Prudential and Moral Reasons

Consider two agents: the pure egoist and the pure altruist. The pure egoist cares only for herself; she is motivated *only* by prudential reasons. The pure altruist cares only for others; she is motivated *only* by moral reasons.¹⁵ As a result, these two agents are motivated in wholly different ways and since motivating reasons characterize our motives, these two agents never act the same way for the same reason, given the schism between their motives.

If we add the orthodox supposition that normative reasons are facts, the scenario implies that these two agents never perform an action on the basis of the same fact. But that's clearly false. For example, each can be moved to save a child because the child is drowning. This is possible because the altruist's and the egoist's goals differ. The altruist is moved, we may suppose, by her goal of respecting and preserving human life. By contrast, the egoist is moved, we may suppose, by her goal of being esteemed and rewarded for saving the child. Nevertheless, it is possible for the same fact, and only that fact, to move each to perform the same act. We must, therefore, reject one of the foregoing claims. As I'll argue, the least plausible claim is the orthodox view that reasons are facts.

Like the case of the Cubs fan and the Cubs hater, this case suggests only that acting on the basis of the same fact does not suffice for acting for the same reason; it does not entail that normative reasons are primary reasons. However, primary reasons offer an enticing explanation of exactly how reasons are more finely grained than facts for it offers an explanation of how reasons resemble facts enough to be confused for them. For example, if normative reasons are primary reasons, then distinct normative reasons, like the egoist's prudential reason and the altruist's moral reason, can overlap on the same fact while nevertheless remaining distinct in virtue of containing different goals.

Indeed, a now-familiar reason exists for thinking that normative reasons contain both goals and facts. Strikingly, Davidson's observation that we attribute motivating reasons in two different ways, which correspond to primary reasons' two parts, also holds for the attributions of normative reasons. For example, we can claim that a *moral* reason to save the drowning child is to respect and preserve human life. Likewise, we can claim that a *prudential* reason to save the drowning child is

¹⁴ Indeed, some like Smith (1994), *op. cit.*, and Susanne Mantel in *Determined by Reasons* (Routledge, 2018) deny that normative reasons can be motivating reasons. The burden on such views is to show why, *pace* common sense, we cannot straightforwardly act for good normative reasons. For this reason, I bracket such views for the remainder of the paper.

¹⁵ For simplicity, assume that these agents are never mistaken about whether a reason is moral or prudential.

to be esteemed and rewarded. If taken at face value, these *to*-clause reasons attributions track differences between moral and prudential reasons, just as more familiar *that*-clause reasons attributions are thought to.

Because these claims involve the attribution of prudential and moral reasons, they provide strong evidence that, for example, the goal of respecting and preserving human life is part of a moral reason while the goal of being esteemed and rewarded is not. However, despite my great sympathy for these claims, I stress that I am simply assuming that they are true in order to illustrate the dual aspect view. Indeed, I am not defending any particular view about which reasons are normative, much less about which reasons are moral or prudential, in particular. I am simply defending a particular model or ontology of reasons according to which they are pairs of facts and goals, rather than just facts or fact-like entities.

In contrast to §2's argument, this argument does not rely on the premise that all normative reasons are motivating reasons, so one cannot resist the conclusion that normative reasons are primary reasons by denying that all normative reasons are motivating reasons. Nevertheless, like §2's argument, it rests on the assumption that the egoist and the altruist act for different reasons despite acting on the basis of the same fact. Given the deep divide in their motives, this claim is compelling. Nevertheless, some may be tempted to insist that the altruist and egoist act for the same reason by insisting that some moral and prudential reasons for an agent to do something are identical. I'll now provide a third argument, which shows that moral and prudential reasons for an agent to act in some way *must* be distinct, so we should accept that normative reasons are primary reasons.

4. Weighing Explanations

Defenders of the orthodoxy may read the case of the egoist and altruist differently. They may take it to show that if a single fact gives both a prudential reason for a certain agent to save the child and a moral reason to do the same, then some facts, like that a child is drowning, are sometimes *both* prudential and moral reasons for the agent to save the child, simultaneously. This implies that, counterintuitively, the egoist and the altruist can act for the same reason precisely *because* they can act on the basis of the same fact. However, as I'll now argue, different kinds of normative reasons for action, like prudential and moral reasons, *must* be numerically distinct if reasons are to play their most central role in moral philosophy, namely, their role in *weighing* explanations of normative facts. In other words, I'll show that weighing explanations require primary reasons' fineness of grain.

When deliberating about what to do, we 'weigh' considerations. For example, that wine is tasty favours drinking another glass. But that I'll feel terrible in the morning if I do favours abstaining. When deliberating about whether to have another glass, I assess how strongly each consideration bears on what to do -- how *weighty* each consideration is -- and act on the consideration I deem weightier. Weighing explanations of what you should do resemble this picture of deliberation. What one should do is explained by the balance of reasons for or against the relevant options.

Reasons are more or less weighty as a function of how strongly each bears on the question of what to do.

Questions about what someone should do can be answered in a similar manner. Of course, 'should' has different senses, which are explained by different groups of reasons. It is controversial whether there is an unrestricted, "all-things-considered" sense of 'should', where all reasons are relevant.¹⁶ But there are clearly senses of 'should' -- like its rational, moral, or prudential senses -- where only certain reasons are relevant. For example, that I've promised to show up on time for dinner tonight is a moral reason to show up on time. But if, on my way to dinner, I see a cyclist get into an accident, then that's a moral reason not to show up on time, and to stay behind to help the cyclist instead. What I should do in this circumstance, morally speaking, depends on the weights of these two reasons. If the promise is especially important -- say, a promise to my beloved to show up on time to our wedding reception -- then that's a weighty reason to show up on time. And if the cyclist's injuries are especially grave, then that's an especially weighty reason to stay behind.

Suppose that, in this particular circumstance, both reasons are equally weighty. In that case, you may keep the promise or you may stay behind. Either option is morally permissible. Crucially, only moral reasons figure in this balance.¹⁷ For example, imagine that the cyclist offers you ten dollars to help them. The reason given by this reward, suppose, shifts the balance of *all* reasons, making the otherwise perfectly counterbalanced set of reasons to help the cyclist weightier than the set of reasons to keep the promise. In that case, you should, all things considered, help the cyclist rather than keep the promise. Nevertheless, even if the cyclist offers a reward to break the promise, you may, morally speaking, keep it. The reward is irrelevant to what you should do, morally speaking, because it gives a merely prudential reason to help, which must be ignored when determining the balance of moral reasons.

More generally, the balance of a set of reasons explains what someone should do, in any of the restricted senses of 'should', 'may', and so on, only if that set meets two criteria. It must include *only* reasons relevant to the restricted sense in question. Otherwise it risks assigning too much weight to one of the options. Likewise, the set of reasons must also include *all* the relevant reasons. Otherwise it risks assigning too little weight to one of the options. When some fact gives two reasons, one relevant to a weighing explanation and one irrelevant, we must include the former and exclude the latter.

The trouble with claiming that a single fact is both a moral and a prudential reason now emerges. Suppose that two children, qualitatively identical in all morally relevant respects, are drowning and only one can be saved. Consequently, you may save either, but you must save one. Moreover, one child is from a rich family; the other child is not. The person who saves the first child will be richly rewarded. Not so for the second. Supposing that reasons are facts, the fact that the rich child is

¹⁶ See, for example, Derek Baker's "Skepticism about Ought Simpliciter" in *Oxford Studies in Metaethics* 13 (2018): 230-252.

¹⁷ I am bracketing the complexities engendered by supererogation for the moment.

drowning is a moral reason to save them. It's also a prudential reason to save them. This creates a dilemma. Either we include that fact when determining the balance of moral reasons, or we do not. If we do include it, we'll have allowed a prudential reason to distort the balance of moral reasons. The prudential dimension of that fact distorts the balance of moral reasons in precisely the same way that the cyclist's reward does, creating a specious moral obligation to save the rich child where there is none. However, if we exclude the fact, then we'll have ignored a morally relevant reason, thereby distorting the balance of moral reasons; we'll falsely suppose that the set of moral reasons includes only the fact that the second child is drowning. Since any reason trivially outweighs an absence of countervailing reasons, we'll be forced to falsely conclude that we must save the second child and may not save the first. Consequently, the assumption that reasons are facts leads to insoluble dilemmas like these where we must either include an irrelevant reason or exclude a relevant one in a weighing explanation.

By contrast, if normative reasons are primary reasons, then we can distinguish the two reasons given by a single fact by associating them with different goals. Perhaps the moral reason given by the fact that a child is drowning is associated with the moral goal of preserving and respecting human life; perhaps the prudential reason is associated with the prudential goal of earning rewards and praise. So long as moral and prudential goals do not overlap, we'll have solved the problem. As a result, we cannot resist the argument for primary reasons as normative reasons given by the case of the egoist and the altruist by asserting that because they act on the basis of the same fact, the two agents act for the same reason. The two reasons given by that fact, one moral and one prudential, *must* be distinct if we are to preserve reasons' most central role in moral philosophy.

I've just argued that weighing explanations depend on distinctions that facts, taken as reasons, don't provide and I've proposed that we should instead weigh more complex objects that resemble primary reasons. I'll briefly consider three responses to this final argument before concluding. The first response concedes that the problem for weighing explanations that I describe is genuine. But, according to this response, it doesn't show that reasons are not facts. All it shows is that we need to make our weighing explanations more complex, somehow.

A first attempt at doing so starts with my claim facts don't provide all the distinctions that a weighing explanation requires. On its face, this claim is odd. Facts can draw arbitrarily fine-grained distinctions. After all, conjunction is recursively iterable. How then could reasons draw more distinctions than facts? Indeed, the problem for weighing explanations above appears to dissolve if we weigh only more finely-grained conjunctive facts rather than the atomic facts initially used to generate the problem. Consider the difference between the following:

1. That child is drowning.
2. That child is drowning and I'll be rewarded if I save them.

For example, it's natural to think that while (1) gives a moral reason to save them, (2) doesn't; it's a purely prudential reason. This latter fact favours saving the child in only a prudential way, and not a moral way, so it does not undermine the weighing explanation of what you should, morally, do.

However, observations about (2) do not show what must be shown to rebut the argument. That argument concerns (1), and its problematic dual significance, not (2). So unless it can be shown through additional argument that facts about (2) imply that (1) is not problematic, discussion of (2) rather than (1) is merely a distraction. Let me also restate the dialectical import of my claim that (1) gives both a prudential and a moral reason. I am not trying to defend a particular view about which reasons are moral; I am simply trying to defend a model of normative reasons inspired by primary reasons. Consequently, though I think it's extremely plausible that (1) gives both a moral and a prudential reason in the context I've described, even if it doesn't, it suffices that (1) makes it extremely likely that *some* facts give both moral and prudential reasons, and that those facts will create similar problems. That likelihood shines a flattering light on the dual aspect view.

On a different way of looking at the problem, it shows only that the *reason relation*, the relational property exemplification of which distinguishes reason-giving facts from more prosaic ones, is more complex than it is ordinarily thought to be. From this point of view, making the reason relation more complex makes weighing explanations correspondingly more complex. For example, we might think that the fact that the rich child is drowning stands in *two* reason relations. It stands in the moral reason relation; as a result, it gives a moral reason to save the child. But it also stands in the distinct prudential reason relation; as a result, it also gives a prudential reason to save the child. Moreover, this account is consistent with the orthodox view for it is consistent with the claim that each normative reason is identical to the fact that gives it.

However, positing complex reason relations solves nothing, all by itself. When we make relations more finely-grained, we make the conditions under which an object satisfies the corresponding relational property more finely grained. For example, when we shift from a conception of an intension that relates sentences to worlds to one that relates sentences to worlds and times, we are able to characterize how the passage of time affects a sentence's truth. But we don't make sentences themselves more finely grained; only the conditions under which they're true. Likewise, when we make the reason relation more finely grained by adding a parameter reflecting morality or prudence, we gain the ability to distinguish how, loosely speaking, morality or prudence make a fact reason-giving. But we do not make reasons themselves more finely grained.

Crucially, to solve the problem, we need to make reasons themselves more finely grained, not simply the conditions under which they are reasons, more finely grained. The problem shows that when offering a weighing explanation of a normative property like the one denoted by moral 'should', the objects weighed in that explanation cannot straddle the distinction between morality and prudence. Distinguishing the moral reason relation from the prudential reason relation doesn't all by itself make weighing explanations more complex for the simple fact that weighing explanations weigh reasons, not reason relations.

Rather, only if we also make the weighing explanation more complex, so that it is sensitive to the parameter in the reason relation that allegedly distinguishes morality from prudence, can this strategy address the problem. There are at least three general worries with this approach, in

addition to the idiosyncratic worries that come with different choices for the parameter in question. First, it is entirely *ad hoc*, motivated by a need to solve the problem and not by independent considerations. Second, because weighing explanations are more complicated than ordinarily thought, on this view, they are less natural, and so less appealing. We cannot simply say that morality requires the act favoured by the balance of moral reasons. We must say that morality requires the act favoured by the balance of moral reasons, *relative to p*, where *p* is the parameter that distinguishes morality from prudence. The naturalness of weighing explanations like the first largely accounts normative reasons' centrality in contemporary moral philosophy. The second, more cumbersome analysis undercuts that appeal, so it offers an only prrhythic solution to the problem. Finally, opting for this solution jettisons our easy account of *to*-clause reason attributions, which strongly suggests locating goals in the reasons themselves and which, as a side-benefit, appears to distinguish morality and prudence thereby solving the problem described above in an independently motivated fashion.

The final response that I wish to raise on behalf of the orthodoxy is the doubt that goals distinguish morality from prudence, despite appearances. If some goals are both prudential and moral, then the dual aspect view does not in fact help to distinguish reasons that must be kept distinct. This worry is genuine. Indeed, depending on how we understand what it is for a goal to be moral or prudential, overlap in goals seems possible. For example, suppose that a goal is moral just when its pursuit tends to make an act morally right. Likewise, suppose that a goal is prudential just when its pursuit tends to make an act prudent. If a view like classical utilitarianism is true, then this account of what makes goals moral or prudential implies that some goals, like that of experiencing happiness, are both moral and prudential. On these assumptions, that goal is both moral and prudential because, other things equal, pursuing one's happiness tends to make an act both prudent and moral.

As a result, proponents of the dual aspect account can't simply deny that a goal is moral just when and because its pursuit tends to make an act right. Rather, they must reject *any* account that similarly predicts overlap in prudential and moral goals. However, there is independent reason for doing so. For example, Christine Swanton argues, roughly, that one acts virtuously just when and because one's act aims at and realizes a virtuous end.¹⁸ According to her, moral goals aren't simply the ones whose pursuit tends towards morally right action; they're the ones that constitute virtuous motives -- whose successful pursuit suffices for virtue.

Swanton's account implies that prudential and moral goals do not overlap. Were overlap possible, a purely selfish agent could act just as virtuously as a morally well-intentioned one, were the first to successfully pursue such a goal. But that's impossible: only the well-intentioned agent acts virtuously. So Swanton's account implies that no virtuous ends are prudential ones. Consequently,

¹⁸ See "A Virtue Ethical Account of Right Action" in *Ethics* 112 (2001): 32-52.; on a similar note, see also Sukaina Hirji's "What's Aristotelian about neo-Aristotelian Virtue Ethics?" *Philosophy and Phenomenological Research* (forthcoming).

the worry that prudential and moral goals overlap rests on assumptions that are contestable on independent grounds.

To be clear, this manner of addressing the third worry does not compromise the dual aspect account's neutrality on the substantive first-order question of which reasons are moral or of which goals are parts of moral reasons. Rather, it only requires committing to a structural claim about whether prudential and moral goals can overlap. Arguments from Swanton and others are grounds for thinking that they do not. But denying that prudential and moral goals overlap does not require affirming that certain goals or reasons are moral or prudential. However, I recognize that more may need to be said to fully dispel this concern.

5. Conclusion

Moral philosophers have largely ignored primary reasons on two grounds. The first is the assumption that normative reasons and motivating reasons are the same kind of things and the dogma that normative reasons are facts. The second is anti-psychologism about all reasons. I've argued that the dogma appears inconsistent with weighing explanations and I've shown how the most important features of primary reasons as normative reasons, their fineness of grain and their explanation of the different ways that we attribute reasons, do not depend on psychologism about reasons. My hope is that these arguments motivate renewed interest in Davidson's primary reasons.