

Eliminating Prudential Reasons

Alex Worsnip

Forthcoming in *Oxford Studies in Normative Ethics*, vol. 8

Penultimate draft; please cite published version where possible

What kinds of reasons are there? There are at least two ways to hear this question. On one way of hearing the question, its answer might go on and on virtually *ad infinitum*. Considerations about many different things can constitute a reason to do something. That a movie has received good reviews is (normally) a reason to go to see it, and this consideration has to do with movie-review-quality, so in this sense movie-review-quality considerations are a kind of reason: the reasons that have to do with movie-review-quality. Likewise, on this reading of the question of what kinds of reasons there are, there will be beer-hoppiness-reasons, how-many-carats-of-gold-reasons, how-strong-were-the-candidate's-letters-of-recommendation-reasons, how-much-time-is-left-before-the-game-starts-reasons, and so on, and so on.

But there is also a different way of hearing that question about what kinds of reasons there are where the answer will be less long-winded, and more philosophically interesting. On this way of hearing it, it is about what kinds of considerations constitute *fundamental* sources of reasons. Presumably, with all of the above examples, there is always some further explanation of why the consideration in question is a reason. That the beer is hoppy is a reason to drink it (or not to drink it) because you *enjoy* (or don't enjoy) hoppy beers, for example. Beer-hoppiness is not a *fundamental* source of reasons for action; considerations of beer-hoppiness don't on their own, in and of themselves, give you reasons to do anything. Beer-hoppiness is not, in Christine Korsgaard's famous phrase, a "source of normativity".

I assume that when philosophers ask what kind of reasons there are, they are asking the second question rather than the first. That is, they mean to categorize kinds of reasons according to the fundamental sources of those reasons.¹ If they were asking the first question, they would need to give a much longer answer than any of the ones they typically give. Given this, it will never be enough to defend putting some kind of reason on the list of "kind of reasons" that considerations of that sort merely can constitute reasons, in the same way that considerations of beer-hoppiness can constitute reasons. To make it onto the list, a category will have to be a *fundamental* kind or source of reasons.

Kinds of reasons that are often put on the list include at least the following: moral reasons, prudential reasons (i.e. reasons given by the acting agent's self-interest), reasons given by the agent's desires (or aims, or projects, or other conative states), aesthetic reasons, epistemic reasons, and reasons of rational coherence. For each putative kind of reason, one can ask whether it really does belong on the list: that is, whether it really is a fundamental source or kind of reason. To use another phrase of Korsgaard's, we can ask "the normative question" about that category.

For helpful comments and/or discussions, I'm grateful to Matt Bedke, Nic Bommarito, Janice Dowell, Daniel Fogal, Paul Hurley, Shelly Kagan, Errol Lord, Ram Neta, Samuel Scheffler, Daniel Wodak, Susan Wolf, a helpful anonymous referee for *OSNE*, and audiences at NYU and the Arizona Workshop in Normative Ethics.

¹ Cf. Smith (1994: 95-6); Crisp (2006: 8).

There are extensive debates about the answer to the normative question for many of the putative kinds of reasons. Thus, some doubt whether morality is a fundamental source of reasons.² Others doubt whether an agent's desires (or other conative mental states) are a fundamental source of reasons.³ Others doubt whether rational requirements of coherence are a fundamental source of reasons.⁴ Yet others doubt whether there are any fundamentally epistemic reasons.⁵

Yet, perhaps surprisingly, prudential reasons (i.e., reasons to do with the acting agent's self-interest) – which frequently appear on lists of different kinds of reasons – seem often to receive a free pass.⁶ That's not to say that the literature is full of clear defenses from first principles of why there *are* fundamentally prudential reasons. But nevertheless, this claim is frequently assumed.⁷ Moreover, the supposition of the unproblematic status of prudence as a source of normativity is revealed by prominent attempts to answer the normative question about *morality* using prudence as a model, arguing that, since we all accept that prudence is a fundamental source of normativity, we should accept the same thing about morality on grounds that are in some respect analogous.⁸ Others go further and treat moral normativity as parasitic on or grounded in prudential normativity.⁹ These styles of argument have dialectical power precisely because it does seem like ordinary folk (and not just philosophers) take the normativity of prudence for granted. It is not a pure coincidence that 'why be moral?' appears as a topic on syllabi for introductory classes, whereas 'why be prudent?' doesn't.

Thus, even if the view that there are fundamentally prudential reasons is not frequently *defended*, I believe that it is frequently *assumed*. This paper challenges this assumption. I will argue that any genuine reason to promote one's own well-being is fundamentally either a moral reason or a reason grounded in one's conative states. That is, there is no non-moral, aim-(or-other-conative-state)-independent reason to promote one's own self-interest. As well as denying that there are *fundamentally* prudential reasons, this also involves a denial that there are *distinctively* prudential reasons: when your own well-being matters, it matters in virtue of having some more general property (such as mattering morally, or being cared about) that other people's well-being can have too. Apart from the contingent fact that we generally care about it more, our own well-being enjoys no special, additional weight that other people's well-being lacks.

1. Preliminaries and clarifications

² Foot (1972) is the paradigm example. Korsgaard's original "normative question" (1996) was about the moral case.

³ Dancy (2000: ch. 2), Scanlon (2004) and Parfit (2011: chs. 2-4) are prominent examples.

⁴ Kolodny (2005) is the paradigm example.

⁵ Rinard (2015).

⁶ Possible exceptions include Rachels & Alter (2005) and Setiya (2015). One can also read Hubin (1980) and Korsgaard (2008: ch. 2) as denying that there are fundamentally prudential reasons, but (unlike me) they both deny that there are *any* reasons that are not grounded in our conative states. Similarly, one can read Singer & de Lazari-Radek (2012) as denying that there are prudential reasons, but (unlike me) they deny that there are any *non-moral* practical reasons. Parfit (1984: 129) comes close at one point to denying that there are fundamentally prudential reasons, but goes on to argue primarily against the logically much stronger view that *all* reasons for action are prudential. In Parfit's later work, he clearly thinks there are prudential reasons; see the next footnote.

⁷ See e.g. Nagel (1970: 16); Crisp (2006: *passim*, esp. 36, 73); Parfit (2011: 40, ch. 6); Broome (2013: 126).

⁸ See esp. Nagel (1970). Parfit (1984: 95, 126, 137-44), inspired by Sidgwick (1907: 418-9), pursues a similar line by arguing for the contrapositive: *if* we do not accept morality as a fundamental source of reasons, we should likewise deny that prudence is a fundamental source of reasons.

⁹ See, e.g., Gauthier (1986).

As I've already said, I'm not denying either that one can have reason to promote one's own well-being, nor that one's own well-being can be a reason to do things (albeit only by mattering morally or by being something one cares about). Additionally, there is a use of 'prudentially ought' whereby to say that one "prudentially ought" to Φ is merely to restate the claim that Φ -ing would all-things-considered promote one's self-interest. Perhaps there is a similar use of 'prudential reason to Φ ' that merely restates the claim that there is some way in which Φ -ing would *pro tanto* promote one's self-interest. Obviously, I am not denying that actions sometimes *pro tanto* promote one's self-interest, and so I cannot be denying that there are prudential reasons in this sense. Similarly, I do not deny that there are considerations that *would* be reasons, hypothetically, if the prudential point of view had authority, or that *seem* to be reasons when one takes that point of view up.

I have said that it is widely assumed that there are fundamentally prudential reasons. But there is one group of philosophers who join me in denying this: namely those who think that *all* reasons are fundamentally given by the acting agent's desires (or other conative states). In my terms, these philosophers think there is only *one* fundamental kind of reason: those given by our desires. There are no other fundamental kinds of reason. *A fortiori*, these philosophers deny that there are fundamentally prudential reasons.¹⁰ However, unlike these philosophers, I wish to argue against the existence of fundamentally prudential reasons *specifically* – not on grounds that generalize to any non-desire-given reason, and in particular not to fundamentally moral reasons. Contra the established philosophical wisdom, then, the normativity of prudence is at least in one way *more* problematic than the normativity of morality.¹¹

It should be obvious from the above that I am using 'prudence' to refer to self-interest, or one's own well-being, and not to the satisfaction of one's desires, or aims, or projects.¹² So when I deny that there are prudential reasons, I am not denying that there are reasons given by these conative states. In fact, as will become important, I *do* think that there are such reasons. At this point, one might ask: what about theories of well-being where well-being simply consists in the satisfaction of one's desires (or preferences)? It looks like the distinction between prudential reasons (as I use the term) and desire-given reasons simply evaporates on such a view.

Note that the only view of well-being that even possibly collapses these two putative kinds of reasons is an *unrestricted* desire-satisfaction theory of well-being, where the satisfaction of *any* desire on an agent's part constitutes an improvement in well-being.¹³ This is a highly implausible view of well-

¹⁰ Of course, someone who thinks that all reasons are grounded in desires can admit that, if we desire or to be moral or to be prudent, we will have reasons to do so; if they like, they can classify reasons given by our moral desires as 'moral reasons', and reasons given by our prudential desires as 'prudential reasons'. But they could also classify the reasons given by our desires for hoppy beer as 'beer-hoppiness reasons'.

¹¹ Here I echo Darwall (2016).

¹² For some, 'prudence' might refer to the latter or be indeterminate between the two. The ordinary meaning of the term isn't of importance to me; I am happy for my usage to be stipulative. Even given my stipulation, I have many opponents, as illustrated above.

¹³ Even then, it's not clear that they collapse entirely. Even if they are extensionally equivalent, we can still ask the question of which one is more explanatorily fundamental in accounting for our reasons. Moreover, if desire-given reasons make reference to *present* desires, whereas the desire-satisfaction theory of well-being makes reference to the satisfaction of *desires over a lifetime* – and these are arguably the most natural ways to develop each view – then desire-given reasons are not equivalent to prudential reasons, even given a desire-satisfaction theory of well-being (see Parfit 1984: 128).

being. I will briefly rehearse two of the most obvious (but also serious) objections to it. First, it is extremely implausible that when some desire of yours is fulfilled, but you believe it not to have been fulfilled, your well-being thereby improves.¹⁴ For example, if you desire that your long-lost son be flourishing, but you believe him not to be flourishing, it is implausible that whether he in fact *is* flourishing affects your well-being. This is not, it bears stressing, to say that it is implausible that you have a *reason* to promote his flourishing, if some means of doing so arises, even if you cannot ascertain the result. On the contrary, I think it is plausible to say that you *do* have such a reason in such a case. That this reason cannot plausibly be construed in terms of a benefit to you reinforces how some desire-given reasons cannot be captured in terms of well-being.

Secondly, an unrestricted desire-satisfaction theory of well-being makes it impossible to deliberately sacrifice your own well-being, except by acting against your own (greater) desire.¹⁵ For suppose that I (most) desire to sacrifice my own well-being for the sake of some greater cause. In such a case, according to the unrestricted desire-satisfaction theory, if the desire to promote the greater cause is satisfied, I thereby become better off. So, if I succeed in satisfying this desire, I have not sacrificed my own well-being after all. This is a conceptually impoverished view. *Pace* first-year undergraduates, it is not impossible for me to deliberately sacrifice my own well-being.

2. A brief defense of aim-given reasons

As I said, I am sympathetic to the claim that there are reasons given by conative states. But I'm not sympathetic to the claim that *all* reasons are given by these conative states. This makes me a "hybrid theorist" in that I accept both reasons that are dependent on our conative states and reasons that are independent of them.¹⁶ (I am most sympathetic to thinking of some of our reasons as given by our *aims*, rather than our *desires*,¹⁷ so I will henceforth talk of 'aim-given reasons'. But much of what I say could be adjusted to make desire or some other conative state the relevant source of reasons.)

It is far beyond the scope of this paper to mount a comprehensive defense of the existence of aim-given reasons. But since I will be relying on aim-given reasons to pick up *some* of the slack upon the elimination of prudential reasons, I should say something about why it's independently plausible that there are aim-given reasons, and that these reasons could not be captured by a framework that contained prudential reasons but lacked aim-given reasons. This will help to lay some groundwork for my argument against prudential reasons.

The cases that best motivate the existence of aim-given reasons are those where we have reasons connected with our projects and relationships, but that are not plausibly classified as either moral or self-interested. On the projects side, plausible cases might include a reason to complete one's novel before one dies, or a reason to contribute to the preservation of the buildings of Venice (when one cares greatly about these things).¹⁸ On the relationships side, plausible cases might include a reason

¹⁴ Cf. Parfit (1984: 494).

¹⁵ Cf. Parfit (1984: 127-9); also Hubin (1980: 75).

¹⁶ See e.g. Chang (2013) and Behrends (2015).

¹⁷ Desires are *pro tanto*, frequently unreflective/unendorsed, and less under our control than aims.

¹⁸ The latter example is from Parfit (1984: 151-2). It is somewhat puzzling (and disappointing) that, having made an impressive case in *Reasons & Persons* (see esp. 94, 133) that such reasons are given by our aims and projects, Parfit's later

to visit a comatose relative in hospital, or to walk for twenty minutes through howling wind and rain to buy one's daughter a hot chocolate (when one cares, again, about the relationships in question).

For any of these examples, there are potential diagnoses that do not treat such reasons as *fundamentally* aim-given. Such alternative diagnoses might take either of two main forms. The first strategy is to argue that these reasons are in fact, at bottom, either moral or prudential. The second strategy is to argue that they need not be moral or prudential, but instead belong to some further, *sui generis* category of reason that is nevertheless independent of our conative states.

The first strategy is, in my view, simply not extensionally adequate. Take the example of walking through howling wind and rain to buy my daughter a hot chocolate.¹⁹ Seeing the happy look on my daughter's face when I give her the hot chocolate is of *some* pro tanto benefit to me, but it might not make me better off *on balance* given the miserable experience of walking through the howling wind and rain. (We can also construct examples where I never *discover* the success of my action benefitting my child, and here it is even more obvious that my well-being need not improve on balance.) Further, it's highly implausible that I'm *morally* required to perform this action. It certainly isn't the thing that I could do with my time and money that would be optimal for the world from an impartial point of view. Nor is it plausible, even if there are special duties of care to one's children, that this act is one of them. But all of this notwithstanding, if I love my daughter and simply *aim* at her happiness, my choice to perform this act seems reasonable. So it seems that what is making the difference here is the aim I have.

Another point is worth making here. Even if one could cook up post-hoc justifications for each of these actions on moral and prudential grounds, it is a crude simplification of human psychology to say that these moral or prudential grounds are always, or even usually, the reasons *for which* we perform such actions. Now, to be sure, there is a distinction between normative reasons and motivating reasons, and it is the former that we are ultimately discussing here. But to the extent that the agents in these examples seem not just to perform actions that *can* be justified in some post-hoc manner, but actually to act *for* good normative reasons, and to deserve whatever credit arises from that, our account should endorse these agents' motivating reasons as genuine normative reasons. So this is another barrier to accounting for these reasons as moral or prudential.²⁰

Turn now to the second strategy, on which the reasons are *sui generis* (but independent of our conative states). On this view, in the case of buying my daughter a hot chocolate, I have a special, *sui generis* kind of (non-moral) reason to promote the happiness of those that I am close to. I have two replies to this strategy. First, the sorts of examples that are connected with our projects and relationships are quite diverse. So, to generalize this strategy appropriately, we will have to introduce many different kinds of *sui generis* reasons: not just reasons given by close ties, but reasons given by the value of preserving beautiful things, reasons given by the value of creating works of art, and so on for all the valuable projects we might have. We will also need some account of how all these different

work (2011: chs. 2-4) steadfastly denies that this can happen, and tries to reduce all such reasons to ones independent of our conative states.

¹⁹ For an 18th century discussion of how acts performed out of parental love need not be self-interested in either their motivation or their normative justification, see Hutcheson (1725/2004: 112-3). For a more recent defense of the view that acts need be neither self-interested *nor* moral in their motivation or justification, see Wolf (2016: 254-5).

²⁰ See Hutcheson (1725/2004: 93-94).

reasons weigh against one another. This is not an obviously fatal objection, but it is striking how much more parsimonious and explanatory an account that unifies all of these reasons, and accounts for their weight, in terms of the *aims* that give rise to them is.

The second reply appeals to the fact that, even if these acts do all in some sense promote something valuable, it does not seem like our reason remains (as strongly) in force irrespective of whether we aim at promoting this value. Someone who cares greatly about preserving the buildings of Venice clearly seems to have more reason to preserve them than someone who does not. Even in the case of the daughter and the hot chocolate, someone who decides that on this occasion she cares more about avoiding the howling wind and rain than getting her daughter the hot chocolate does not seem unreasonable. But equally, so does the person who decides that he cares more about getting his daughter the hot chocolate. This isn't plausibly because the value of each option is *precisely tied*. In fact, it's not even quite the case that both choices are reasonable, *simpliciter*. It *would* seem like some kind of mistake for the person who cares more about staying dry to buy the hot chocolate, or for the person who cares more about buying the hot chocolate to stay put (unless, in doing so, they thereby constitutively revise what they care about most). So, it seems, there's an element of latitude here for agents to affect their own reasons by shaping their own aims and cares.²¹

A final option for the denier of aim-given reasons is simply to stick to his guns and deny that I have any reasons in these cases, above and beyond the moral value of these acts and their contribution to my own well-being. I find this intuitively implausible. But moreover, this is a particularly tough position for someone who *admits* fundamentally prudential reasons to take. Such a person will be forced to defend the view that, while I may (or even must) give my own well-being additional weight on top of whatever significance it has morally, I may not give my daughter's well-being additional weight on top of whatever significance it has morally. This strikes me as an ugly consequence.²² So those who admit prudential reasons should find it particularly hard to deny aim-given reasons. I respect the pure moralist who thinks that moral reasons are all there are, and that aims don't make a difference. But dialectically, in arguing against prudential reasons, I can ignore this view.

3. Against prudential reasons I

I will make two, related arguments against (fundamentally) prudential reasons. The first follows on from some of the points I just made in defending aim-given reasons. Anyone who thinks there are fundamentally prudential reasons must hold that your own well-being carries some weight, in addition to whatever weight it has morally, that is not contingent on your conative states. But if this is so, then (at least normally) you are *required*, from the all-things-considered point of view, to assign your well-being this extra weight in your deliberations. In my view, this is implausible. It is one thing to say that you are permitted to give it this weight *if you want to*. It is quite another to say that you make a mistake if you do not.

²¹ Here I follow Chang's (2013) strategy of appealing to aim-given reasons as a way of resolving cases that objective values leave undecided.

²² But see Crisp (2006) for a defense of it.

The implausibility here is starkest if we consider a moral theory that (i) assigns your own well-being equal weight already within the moral domain, and (ii) rarely, if ever, allows that you are in a situation where there are no relevant moral considerations. Utilitarianism is, of course, just such a view.²³ If we combine utilitarianism about morality with additional, fundamentally prudential reasons, then it looks like, all-things-considered, you are *required* to give yourself more weight than others – such that you would always be making a mistake if you did not do so.²⁴ This is a particularly stark consequence.²⁵

But I think the worry generalizes beyond moral theories of this sort. Here it is helpful to make a comparison with the projects and relationships that we considered in the previous section. There, I contended that you have latitude to affect the existence and weight of reasons to do things like preserve Venetian buildings, complete your novel, or buy your daughter a treat. Of course, it's certainly possible that you might have moral reasons to do any of these things, depending on how we tell the story, and different moral theories will say different things about this. But regardless of which of these moral theories we accept, whether you have *additional* reasons to do these things is something that depends on your conative states. Since you have some latitude over these conative states (e.g., the capacity to set aims), you thus have latitude over the existence and weights of these reasons.

My challenge to the defender of prudential reasons is to explain why your own well-being is different from these projects. Again, your well-being may matter morally, and different moral theories will say different things about this. But why should it be *mandatory* to give your own well-being further, non-moral weight, regardless of how much you care about it, if it is not mandatory to give such further non-moral weight to the well-being of others, to those you are close to, or to other things that might be among your projects? This question arises regardless of what moral theory we accept, and how much moral weight it gives to one's own self-interest. My contention is that, just like your projects and relationships, your own well-being is something that gives you *additional*, non-moral reasons only insofar as you care about or aim at it. Often, you will care about it a lot. But if, as a matter of fact, you

²³ One can imagine a non-standard version of utilitarianism that identifies itself not as a moral theory but as a theory of what one all-things-considered ought to do. On such a view, one has moral reasons to impartially promote the well-being of everyone *except* oneself, plus (fundamentally) prudential reasons to promote one's own well-being that are of precisely equal weight to the moral reasons that one has to promote a single individual's well-being, thus adding up to an all-things-considered view according to which one ought to give impartial, equal weight to everyone. This concedes the spirit, if not the letter, of my view. Moreover, while I think there's some appeal to the idea that one's own well-being doesn't count *morally*, I'll argue below that if one finds this persuasive, one should also think that one's own well-being doesn't carry fixed, desire-independent normative weight of some *sui generis* kind either.

²⁴ Crisp (2006: 131-44), for example, holds that one has both "impartial reasons" – which are determined by a broadly utilitarian framework, in which one already counts equally with others – *and* self-interested (prudential) reasons, that give one's own well-being additional weight on top of what it already carries from the impartial perspective. (He does not want to call the impartial reasons "moral reasons", but the structural point remains.) Cf. also Parfit (2011: 136-9).

²⁵ Parfit (2011: 140) tries to avoid this, while also preserving "reasons to be specially concerned about our own well-being", by holding that although we are all-things-considered permitted to give equal weight to ourselves and others, we are also permitted to give more weight to ourselves than others, but *not* to give more weight to others than ourselves. First, I am not sure that this stance really vindicates the existence of (fundamentally) prudential reasons. Although it entails that our well-being *can* enjoy additional weight that the well-being of others cannot have, it concedes that we lack categorical (and thus, conative-state-independent) reasons to give ourselves extra weight. Secondly, I still find Parfit's revised view objectionable. *If* it is permissible to give myself extra weight of this sort, I do not see why it should not also be permissible to give others extra weight.

care more about something else, you make no mistake by discounting it (unless you thereby violate a moral obligation).

At this point, it might be conceded that you have *some* latitude over the normative weight of your own well-being. But, it might be continued, isn't there some minimum baseline below which you may not discount your own well-being? I do not have to deny this. I simply have to insist that if there is such a minimum baseline, it is whatever weight your well-being has morally. This is not just an arbitrary just-so story. A genuine *moral* obligation to take your own well-being into account is just the sort of thing that one could not diminish the significance of by an act of will, or indeed will away entirely.

But what if there is never any moral obligation to take one's own well-being into account?²⁶ In that case, it seems that my appeal to moral considerations as providing the baseline below which one may not discount one's own well-being will not work. For such views, conjoined with a denial of fundamentally prudential reasons, yields the result that one has no reason at all to promote one's own well-being independently of one's conative states. Let me make several points here. To start with, to the extent that one finds this result implausible, it's not clear that it's the assumption that there are no prudential reasons – as opposed to a theory of morality that gives one's own well-being no weight at all – that is to blame. The contention that one's own well-being is has *no* moral weight is a very strong claim. It would be, in itself, a highly interesting result if accepting this was the only way of hanging on to prudential reasons.

Moreover, to the extent that one *is* moved by the thought that one's own well-being never provides one with moral reasons, I think one should also be willing to bite the bullet on the claim that one has no reason to promote one's own well-being independently of one's conative states. The key thought in motivating a radically other-regarding view of morality is that I have a special authority to *discount* or *waive* considerations of my own well-being, in a way that I do not have a special authority to discount or waive considerations of the well-being of others.²⁷ I may simply decide to sacrifice myself, whereas I may not simply decide to sacrifice you. But if it is then contended that, although I do nothing *morally wrong* in sacrificing myself, my act may still be what I *all-things-considered* ought not do, then this authority to discount or waive considerations of my own well-being is, after all, denied to me. This personal authority to discount or waive considerations of my well-being, beyond any moral weight they have, is just what I have been insisting upon. If one also holds that this personal authority is so great that such considerations cannot even have *moral* weight, it is perverse to let them back in as distinctively prudential reasons that I do *not* have the authority to discount. Thus, while I take no stand on the claim that we have full authority to waive all considerations of our own well-being, I do think that it is plausible to think of the truth of that claim as being determined by the answer to the question of whether such reasons can be grounded morally.

As I just said, I am not endorsing the claim that we have full authority to waive all considerations of our own well-being. But I do think there is *something* plausible about it. A more intermediate view might affirm a *partial* asymmetry between the moral weight of one's own well-being

²⁶ For philosophers that endorse this view, see e.g. Ross (1939: 72, 273-4), Singer (1959); Stocker (1976); Finlay (2007: 139-42).

²⁷ See, e.g., Singer (1959: 202).

and the moral weight of the well-being of others. One view that seems especially attractive is that any case where one *does* have a moral duty to oneself has to be one where there is a certain distance between the one's present aims and those of the future self that may be adversely affected. This distance creates a kind of mock-interpersonal or quasi-interpersonal structure that allows for a moral obligation to arise. Here one lacks the authority to waive one's own well-being precisely because one does not speak on behalf of one's future self. Thus, for example, a 20-year old who only cares about present pleasure might have a moral duty to her future self, who will care about living to see her grandchildren grow up, not to engage in risky or life-shortening behavior.²⁸ Conversely, though, if her future self largely shares (or would share) her current aims – living fast and having a significantly briefer but more intensely pleasurable life – then she will lack any obligation to the contrary.

Finally, you may wonder whether the implausible results of accepting fundamentally prudential reasons can be avoided if one *also* accepts aim-given reasons.²⁹ As I have argued at the end of the previous section, accepting aim-given reasons does make the defender of prudential reasons *somewhat* better off. Nevertheless, a view that embraces both prudential and aim-given reasons still says that one is required to give one's own well-being a certain amount of weight, beyond its moral significance. I have been questioning why your own well-being enjoys this mandatory, desire-independent weight, on top of its moral significance, and suggesting that in fact it enjoys no special priority over many other things you might care about (except insofar as it enjoys any moral significance that they do not). Though allowing aim-given reasons allows those other things to matter (nonmorally) too, it still gives self-interest a privileged status by having it matter (nonmorally) *over and above the extent to which you care about, or aim at it*.

This section has appealed to many overlapping considerations. You may want a premise-conclusion argument to take away from it. Here is one such argument:

- (1) If you had any fundamentally prudential reasons, they would be given by your own well-being.
- (2) When your own well-being gives you a reason, this reason either depends on your conative states (i.e., is hypothetical) or doesn't (i.e., is categorical).
- (3) If your own well-being gives you a reason that depends on your conative states (i.e., is hypothetical), then this is not a fundamentally prudential reason.
- (4) The only way that your own well-being can give you a categorical reason is by mattering morally.
- (5) If your own well-being gives you a reason by mattering morally, then this is (again) not a fundamentally prudential reason.

Therefore,

- (6) You do not have any fundamentally prudential reasons.

²⁸ The relevant distance need not always be temporal. Consider Parfit's (2011: 73-82) agony case, where one now lacks any desire not to suffer severe agony in the near future. Assuming that one cannot actually be *suffering* agony without (*at that time*) desiring it to stop, there is very significant difference between the conative states of one's present, agony-indifferent self, and one's near-future, agony-suffering self. It is not implausible to say that one has a moral reason to prevent one's future agony in this case.

²⁹ Chang (2013) and Behrens (2015), both hybrid theorists who think that some reasons are grounded in conative states and some aren't, do not seem to reject prudential reasons.

This rendering of the argument may not be enormously helpful, because all the action is in premise (4). Premise (2) is a tautology, and premises (1), (3) and (5) are all just definitional consequences of what I mean by ‘fundamentally prudential reason’. Premise (4), by contrast, is controversial, and I acknowledge this. By offering it as a premise in this argument, I do not mean to imply that it is *prima facie* self-evident. Rather, I have been trying in this section to support it.

4. Against prudential reasons II

My second argument against prudential reasons is related, but distinct. It concerns the difficulties in the “weighing” of reasons that admitting fundamentally prudential reasons creates.

A picture on which moral reasons and prudential reasons weigh against one another is part of a broader trend away from monism in the theory of normativity. On monistic views, it was assumed that there could only be one fundamental source of normativity. Prudence was thus seen as not as a source of reasons that weigh against other reasons, but rather as supplying a comprehensive theory of what one all-things-considered ought to do – egoism.³⁰ It thus conflicted with impartial morality (conceived also as a monistic theory) in a very strong sense, since only one such monistic theory could be true. By contrast, a pluralist view takes it that there are multiple sources of reasons – morality and prudence being among them. Such reasons may conflict in the weaker sense that they support different courses of action. But there is nothing like a logical contradiction in the possibility of simultaneous, conflicting reasons. These different reasons, on the pluralist picture, weigh against one another to supply an answer as to what one all-things-considered ought to do.

If we posit both moral and prudential reasons, we face the question of saying what the relative weights of such reasons are. Here it is hard to see what a principled answer might be, or how we might go about discovering it. Take a very simple choice like that between giving money to charity or going out to a nice dinner tonight. Presumably I have some moral reason to donate to the charity, and (if there are prudential reasons) some prudential reason to go out to dinner. But there is no obvious way to see which reason is stronger, or perhaps even to know what might in principle make it the case that one reason is stronger than the other.³¹ I cannot just answer this question by doing moral theory and asking whether I am *morally* obligated to give the money to charity. For this is not yet to say that whatever moral reasons are provided by this obligation outweigh the prudential ones.

One principled and clear answer is that moral reasons always override prudential reasons. This might not make prudential reasons entirely irrelevant, if there are cases where there are no moral reasons in play. Still, many of those who posit prudential reasons do not want to say that they are

³⁰ A paradigm example of this kind of monistic view is Sidgwick (1907). This may sound surprising, since Sidgwick is famous for his “dualism of practical reason” between impartial utilitarianism and egoism. But Sidgwick’s “dualist” idea is not that both impartial utilitarianism and egoism are sources of *pro tanto* reasons. Rather, it is that there is a fundamental undecidability between impartial utilitarianism and egoism as monistic theories of all-things-considered reasonable conduct. It is precisely because Sidgwick assumes that we *need* a monistic theory that this undecidability is supposed to be so troubling and that Sidgwick concludes, in light of it, that his book has ended in failure. For a recent defense of impartial utilitarian monism, see Singer & de Lazari-Radek (2012).

³¹ This may have been why Sidgwick did not take the *pro tanto* view. As his book’s title suggests, Sidgwick wanted a *method* of ethics, and this stance supplies no such method.

always overridden by moral reasons. Nor do they want to say that the moral reason outweighs the prudential reason simply because the charity donation will bring more well-being to others than the dinner will bring to me. As we saw in the previous section, most advocates of prudential reasons think that one's well-being counts in a distinctive, *extra* way in addition to whatever impartial weight it already has from the moral point of view.

Moreover, even if one were to take a view on which morality overrides prudential reasons, I have already argued in section 2 that moral reasons and prudential reasons (assuming we admit the latter) fail to exhaust the reasons that there are. Everyone, and the defender of prudential reasons in particular, ought to admit the existence of further, aim-given reasons. And then the question of the interaction between prudential reasons and aim-given reasons arises. Here an overridingness claim is simply not plausible. Recalling our case of buying my daughter a hot chocolate, it was a stipulation of the case that overall, my well-being is (slightly) worsened by buying my daughter the hot chocolate. But doing so, when I care about her happiness, is not unreasonable. So we need some account of the relative weight of prudential and aim-given reasons. The task here seems even harder than the task of giving an account of the relative weight of moral and prudential reasons.

An alternative picture is one on which each individual reason just has some brute weight, and there is nothing systematic to say about which kinds of reasons have greater weight; it just depends on the particular reasons involved. I don't have a knockdown objection to this view, but I do think we should avoid positing such brute, interacting weights – especially to reasons of very different kinds – where we can. To put the point somewhat tendentiously, if we can just about make peace with the idea that there are normative principles written in the sky of Plato's heaven, the idea that there are also, written in that same sky, quasi-mathematical 'weights' assigned to every consideration in every conceivable situation that an agent might face, seems too much. If the relative weights can be explained in terms of some common factor, the problem isn't so great. But that's unlikely to be so when the reasons are supposed to be of different *fundamental* kinds.

Now some philosophers who endorse reasons of various different kinds, such as Parfit (2011: 132) and Broome (2013: 61), do admit that there is a certain amount of imprecision when it comes to the weights of reasons. They think that there are some cases where the moral reasons outweigh the prudential ones, other cases where the prudential reasons outweigh the moral ones, and other unclear cases where neither determinately outweighs the other. It seems implausible to say that the cases in this last category are ones where the moral and prudential reasons have *very precise* weights than just happen to be identical, so instead these philosophers say that in these cases there is some indeterminacy.

For several reasons, I don't think this concession of indeterminate cases helps all that much. First, both Parfit and Broome claim that in cases where it's indeterminate whether the total balance of reasons supports some action A or some incompatible action B, it's permissible to perform either action: that is, it's neither the case that one ought to do A or that one ought to do B. This yields a fairly plausible result about cases where it's hard to weigh the moral and the prudential reasons: one may permissibly go either way. But it is not clear to me that Parfit and Broome can really legitimately help themselves to these claims. On both of their theories, what one ought to do is equivalent to what one has most reason to do. If it's genuinely indeterminate whether one has most reason to do A, then,

it should also be indeterminate whether one ought to do A (and likewise for B). This differs from the result that, determinately, it's neither the case that one ought to A nor that one ought to do B, and that both are permitted.

Secondly, it's important to distinguish between cases where it's *not intuitively obvious* which reason is stronger, and cases where there's no *determinate fact of the matter* about which reason is stronger – where the former property of a case is epistemic and the latter metaphysical. If all cases of the former kind are of the latter kind, then the present view yields indeterminate normative verdicts about all non-obvious cases – which puts very strong constraints on what any normative theory could ever tell us about what we all-things-considered ought to do. If, on the other hand, not all cases of the former kind are of the latter kind, then we are back with the question of how to tell whether the moral reason or the prudential reason is stronger in the cases that are determinate but non-obvious.³²

Thirdly, even in the apparently obvious cases, the present account has no explanation of what *makes* it the case that one reason is stronger than the other. Thus it still countenances brute weights for reasons in a way that seems (to me) objectionable. Admitting indeterminacy in some cases does not change this.

I have been trying to make trouble for views that allow fundamentally prudential reasons, by pointing to difficulties in the weighing of moral and prudential (or prudential and aim-given) reasons. But an obvious objection will surely have occurred to you. Even if we eliminate prudential reasons, we are still left with moral reasons and (by the lights of section 2, at least) aim-given reasons. And, one might think, all the same problems about reason-weighing arise here. So these problems, one might think, do not make trouble for prudential reasons *specifically*. Rather, they make trouble for any pluralist theory of the sources of normativity. We are caught between the rock of monistic theories, which are extreme and extensionally inadequate, and the hard place of pluralistic theories, which are plagued by problems of weighing.

In what remains, I will argue that this is in fact not so. Eliminating prudential reasons allows us to give an attractive and plausible picture of reason-weighing without accepting monism. I can only sketch this picture, and cannot argue for it here.

The picture I like is this. There are certain moral obligations that act as genuine, overriding side-constraints on the permissible pursuit of our aims. Which moral considerations amount to true overriding obligations of this kind is a matter settleable by moral theory; it is not a further question to be established once moral theory has been concluded. It might be a fact of moral theory that no true moral theory would issue overriding obligations that would preclude us from satisfying certain aims in certain ways. In that case, it won't be that our aim-given reasons *outweigh* would-be overriding moral reasons in such cases, but rather that morality makes way for the pursuit of our aims in such cases, by not issuing over-riding obligations.

Does this mean that whenever morality *does* speak, it is overriding? No, for in addition to these over-riding moral side-constraints, there may be morally supererogatory acts. So, what are we to say about the relative weights of supererogatory moral considerations and aim-given reasons? My answer is that *these* relative weights are up to us. Indeed, being supererogatorily moral is, from the standpoint

³² Crisp (2006: 136-7) tells us that we simply have to exercise *judgment* in such cases. But this understates the depth of the challenge, since we are lost as to what *basis* we could make such a judgment on.

of what one has most reason to do, just another possible aim or project, and one will have reason to act on supererogatory moral considerations just to extent that one aims at doing so.³³ In summary, then, the weights of our aim-given reasons – including reasons given by supererogatory moral aims – are determined by the depth, strength, or resilience of those aims – but there are also certain overriding moral considerations that put side-constraints on whether it is permissible to act on these aims, effectively outweighing the reasons that they provide.

This view is not monistic: it allows for both moral and aim-given reasons. Nevertheless, it avoids positing any situation where we are forced to compare reasons of very different kinds without a clear metric for doing so. When the conflict is between an overriding moral obligation and an aim-given reason, the over-riding moral obligation always wins. When the conflict is between a supererogatory moral consideration and a reason given by some other aim, the relative weight of each is determined by whether one more strongly or deeply aims at the supererogatory or at the conflicting aim. Admittedly, it will not always be easy to know whether a moral consideration is overriding or supererogatory. But that is a matter for moral theory, and resolving it poses no special challenge that other difficult moral questions do not pose.

Someone who admits (fundamentally) prudential considerations cannot mimic this picture. Note that there are three (new) interactions that would need to be explained when prudential considerations are introduced:

- (i) The interaction between overriding moral considerations and prudential considerations
- (ii) The interaction between supererogatory moral considerations and prudential considerations
- (iii) The interaction between prudential considerations and aim-given considerations

Case (i) is simple: since the moral considerations in question are overriding, they override the prudential considerations. But there is still the challenge to explain cases (ii) and (iii). Here one might try mimicking my approach to the weighing of supererogatory and aim-given considerations, by saying that in cases (ii) and (iii), the weights are determined by one's aims. But – given that we have already said that prudential considerations are overridden in case (i) – that would be to concede that prudential considerations in principle fail to make a difference in any case where they come into conflict with other considerations, except insofar as one aims at acting on them. And that is effectively to concede that there are no aim-independent prudential reasons.

In this second argument, I have been appealing to the difficulty of giving a principled account of the weights of reasons of fundamentally different kinds. Some philosophers are just not moved by this problem. They think that any account of normativity whatsoever will ultimately have to countenance the weighing of reasons of different kinds, where there are no general principles about the relative weights of these different reasons. For instance, they may think that *within* moral theory, we are forced to countenance conflicts between different moral values, where different considerations must each be assigned a weight, and there is little to be said in general about these relative weights. I myself am not resigned to this outcome; in moral theory, just as in the theory of practical reason more

³³ Of course, acting on supererogatory moral considerations is still more morally praiseworthy than acting on other aims.

broadly, I am inclined toward accounts that minimize the need for brute weights attaching to different, competing considerations. But I recognize that not everyone shares this instinct. In dealing dialectically with those who do not, I may need to confine myself to my first argument. Still, I hope that the foregoing has at minimum shown that there is a picture of the interaction of different kinds (viz. moral and aim-given) reasons that is *prima facie* attractive, that avoids some of the potential complications of weighing different kinds of reasons without recourse to monism, and that the defender of prudential reasons is precluded from accepting.

5. A brief final point

In both of my arguments above, I have appealed to the idea that we have *latitude* over the normative weight of various kinds of considerations: over the weight of supererogatory moral considerations, of prudential considerations, and of other aims, projects and relationships that we have. A final question might be this: can a framework that tries to explain such reasons as fundamentally aim-given explain this latitude? For example, I say that when one *does* aim at one's own well-being, one has reasons of a certain weight to advance one's own well-being. But can't my point about the all-things-considered permissibility of sacrificing one's own well-being for other things still be made even when one *does* aim at one's own well-being?

However, the intuition I have been appealing to is really that when one *decides* to sacrifice one's own well-being because one *cares more* about some other goal, this is all-things-considered permissible. In such a case, the caring and/or deciding plausibly constitutively make it the case that one aims at the other goal more deeply or strongly than one aims at one's own well-being. By contrast, if we stipulate that the case involves acting on the other aim without deciding to do so or caring about it more, the intuition that one is not doing as one all-things-considered ought to is, in my view, much easier to get. Moreover, in a case where you aim more strongly at your own well-being than the other goal, we might still say that it is permissible to act on the other goal, as a way of meaning that it is permissible to *revise* your aims so as to aim more strongly at the other goal, and then act accordingly.

Thus I think that treating certain reasons as aim-given does make sense of our having latitude over their weight.³⁴ We have this latitude because we have the power to set our own aims.³⁵

6. Conclusion

These days, most philosophers reject normative egoism. But if I am right, they have not rejected it deeply enough.³⁶ They have (rightly) rejected the view that *only* my own well-being can provide me with reasons for action, or that I always all-things-considered ought to pursue my own self-interest. But they still cling to the belief that there is *some* way in which my own well-being is prioritized over other things I might aim at; that there is *some* way in which my own well-being provides me with

³⁴ Cf. also Chang (2013: 171-3).

³⁵ But perhaps not our own desires, and this is one reason for preferring an account in terms of aims to an account in terms of desires.

³⁶ See Ross (1939: 273-4) for a similar thought, though he is talking about the rejection of *moral* obligations to oneself.

reasons that is intrinsically not a way in which the well-being of others can provide me with reasons. This idea is deeply engrained in Western culture. Nevertheless, I have argued, we should reject it.

References

- Behrends, J. (2015). "Problems and Solutions for a Hybrid Approach to Grounding Practical Normativity," *Canadian Journal of Philosophy*, 45/2: 159-78.
- Broome, J. (2013). *Rationality Through Reasoning*. Chichester: Wiley-Blackwell.
- Chang, R. (2013). "Grounding Practical Normativity: Going Hybrid," *Philosophical Studies*, 164: 163-187.
- Crisp, R. (2006). *Reasons and the Good*. Oxford: Oxford University Press.
- Dancy, J. (2000). *Practical Reality*. Oxford: Oxford University Press.
- Darwall, S. (2016). "Making the "Hard" Problem of Moral Normativity Easier," in Lord & Maguire (eds.), *Weighing Reasons*. Oxford: Oxford University Press.
- Finlay, S. (2007). "Too Much Morality," in Bloomfield (ed.), *Morality and Self-Interest*. Oxford: Oxford University Press.
- Foot, P. (1972). "Morality as a System of Hypothetical Imperatives," *Philosophical Review*, 81/3: 305-316.
- Gauthier, D. (1986). *Morals by Agreement*. Oxford: Oxford University Press.
- Hubin, D.C. (1980). "Prudential Reasons," *Canadian Journal of Philosophy*, 10/1: 63-81.
- Hutcheson, F. (1725/2004). *An Inquiry into the Original of Our Ideas of Beauty and Virtue*. Indianapolis: Liberty Fund.
- Kolodny, N. (2005). "Why Be Rational?," *Mind*, 114/455: 509-563.
- Korsgaard, C.M. (1996). *The Sources of Normativity*. Cambridge, UK: Cambridge University Press.
- (2008). *The Constitution of Agency*. Oxford: Oxford University Press.
- Nagel, T. (1970). *The Possibility of Altruism*. Princeton: Princeton University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- (2011). *On What Matters: Volume 1*. Oxford: Oxford University Press.
- Rachels, S. & Alter, T. (2005). "Nothing Matters in Survival," *Journal of Ethics*, 9/3-4: 311-330.
- Rinard, S. (2015). "Against the New Evidentialists," *Philosophical Issues*, 25: 208-223.
- Ross, W.D. (1939). *Foundations of Ethics*. London: Oxford University Press.
- Scanlon, T.M. (2004). "Reasons: A Puzzling Duality?," in Wallace, Pettit, Scheffler & Smith (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press.
- Setiya, K. (2015). "Selfish Reasons," *Ergo*, 2: 445-472
- Sidgwick, H. (1907). *The Methods of Ethics*, 7th ed. London: Macmillan.
- Singer, M. (1959). "On Duties to Oneself," *Ethics*, 68/3: 202-205.
- Singer, P. & de Lazari-Radek, K. (2012). "The Objectivity of Ethics and the Unity of Practical Reason," *Ethics*, 123/1: 9-31.
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.
- Stocker, M. (1976). "Agent and Other: Against Ethical Universalism," *Australasian Journal of Philosophy*, 54/3: 206-220.

Wolf, S. (2016). “Meaningfulness: A Third Dimension of the Good Life,” *Foundations of Science*, 21: 253-69.