

Second-personal Reason-giving

Fabienne Peter, f.peter@warwick.ac.uk

Paper prepared for *Grazer Philosophische Studien* (90, 2014: 21 - 34), special issue on Stephen Darwall's work

Penultimate draft; please cite published version.

Abstract

David Enoch has recently objected to Stephen Darwall's account of second-personal reason-giving that the phenomena that Darwall focuses on can be fully explained without resorting to second-personal reasons. In this paper, I shall argue, against Enoch, that second-personal reason-giving matters. My account of second-personal reason-giving differs from Darwall's however, as it accepts that some of the phenomena Darwall focuses on can be reduced to the more standard form of reason-giving.

Introduction

The key insight that Stephen Darwall develops in *The Second-Person Standpoint* is that our interactions with others can be of fundamental normative significance. It is nothing new, of course, that our interactions with others may causally affect us. If you are in a bad mood, I might not enjoy being around you as much as I normally do.

If you tell me about a recent trip you made, I might form a desire to visit those places too. And if you tell me that you have just read that it will rain again this afternoon, this might cause me to form a corresponding belief. But we can establish those causal effects without establishing anything about the normative significance of these interactions. The normative significance of our interactions with others depends on how they relate to our (normative) reasons for action or for attitudes such as beliefs, desires, etc. To keep things simple I shall focus here on reasons for actions – the case of practical normativity.

It is also fairly uncontroversial to say that our interactions with others may be normatively significant in the sense that they trigger certain reasons. For example, if I act dismissively towards you, this may give you a reason to express blame or to avoid me. The reason to blame or to avoid disrespectful people is not, I take it, created by my action. It has been there all along. But my action may trigger this reason, such that your normative situation is now no longer the same as it was before I interacted with you in this way.

Darwall has drawn our attention to the possibility that our interactions with others can be normatively significant in a more fundamental way: they may not just trigger reasons that have been there all along; they may also create reasons that would not exist without this interaction. Second-personal reasons are of this kind.

One of Darwall's favourite examples is the following. Suppose I accidentally stepped on your toe and you claim that I should remove my foot from on top of yours.

Darwall argues that this gives me a reason to remove my foot that was not there before you claimed that I should, a reason that is different from other reasons I might have to remove my foot, say reasons grounded in sympathy. It is a second-personal reason that depends for its existence on properties of our relationship. It depends on whether you have the authority to make such claims and, vice versa, on whether I am accountable to you in this respect. Such second-personal reason-giving, Darwall claims, plays a role in requests, commands, promises, and similar practical scenarios and its role is essential for explaining moral obligations.

David Enoch (2011; 2014) has recently challenged Darwall's account of second-personal reason-giving. Enoch argues that Darwall's account is not only metaphysically fishy – it appears that reasons are created out of thin air – it is also unnecessary: we can account for the kind of normative phenomenon that Darwall focuses on – Enoch calls it “robust” reason-giving – in terms of triggering reason-giving.

My aim in this paper is to defend the normative significance of second-personal reason-giving against Enoch's objections. My defense will, however, offer an interpretation of second-personal reason-giving that differs somewhat from Darwall's own interpretation and that assigns it a different place in our normative geography, as it accepts that some of the phenomena Darwall focuses on can be reduced to the more standard form of reason-giving.

Darwall on Second-Personal Reasons

As Darwall influentially put it, second-personal reasons are reasons whose

“validity depends on presupposed authority and accountability relations between persons and, therefore, on the possibility of the reason’s being addressed person-to-person” Darwall (2006, 8).

Second-personal reason-giving is from within what Darwall calls the second-person standpoint. The second-person standpoint is a web of four concepts – reason, claim, practical authority, and accountability – where each entails the others. Someone’s valid claim gives someone else a reason to act accordingly. To use Darwall’s own example again, your claim that I move my foot from on top of yours, gives me a reason to do so. And it does so, not because your claim triggers my sympathy, but because we mutually acknowledge our authority and accept our accountability to each other in cases like that – we each accept that you have a right to make claims of this sort on me and, vice versa, that I have a right to make this sort of claim on you. In other words, what makes your claim valid is that I accept my accountability to you in this regard or, which is the same thing, that I accept your authority to make such claims. Without that presupposed relationship of authority and accountability, your claim would not be valid and would thus not give me a reason to act accordingly.

The web between the four concepts characterizes the relationship between moral agents and exhaustively explains, according to Darwall, the normative grip of each of the components. What is distinctive about the second-person standpoint is that it locates a source of normativity in the relationship between moral agents – not in

individuals as such (first-person standpoint) and not outside of their relationship (third-person standpoint).

The divisions that Darwall draws between the second- and third-person standpoints, on the one hand, and, on the other, between the second- and first-person standpoints are not equally sharp. The significance of the second-person standpoint in morality is linked to a denial of a third-personal, or fact-relative, source of moral normativity. So the distinction between the second- and the third-personal standpoint in ethics is, on Darwall's account, a sharp one. As Darwall explains (2006, 8): second-personal reasons "simply wouldn't exist but for their role in second-personal address". As I interpret it, the key claim that *The Second-Person Standpoint* makes in this regard is a constructivist one: second-personal reasons are agent-relative in the fundamental (or metaphysical) sense that Korsgaard (1996) has identified: they are reasons that depend on our agency for their existence.

Because second-personal reasons depend on our agency, first-personal considerations are part of the second-person standpoint (Pauer-Studer 2010). But this does not imply that the second-person standpoint reduces to the first-person standpoint. Instead, the second-person standpoint serves to qualify first-personal considerations; it is necessary to establish what we can validly claim from each other.

Enoch on Robust Reason-Giving

Enoch accepts that there seems to be something special about reason-giving in contexts such as requests or commands that a theory of practical normativity needs to be able to explain. But he rejects the explanation that Darwall's account of second-personal reasons offers.

Enoch (2011, 2014) distinguishes between three different senses in which interactions with others might give you reasons for action. The first is the epistemic sense. Epistemic reason-giving occurs, for example, if you draw my attention to a relevant practical consideration I had overlooked. To illustrate, you might point out to me that it is my grandmother's birthday tomorrow. This gives me a reason to call my grandmother. But the reason for action was there all along, independently of our interaction.

Second, there is the reason-triggering sense that I have already mentioned. Triggering reason-giving occurs if the interaction with others manipulates the non-normative circumstances such that a latent reason becomes active. Again, the reason is there all along; it is not created by the interaction. Enoch gives the example of someone setting a foot on the road, thus activating the reason for drivers "to-stop-should-a-pedestrian-start-crossing" (2011, 4).

Enoch calls the third sense "robust reason-giving". That is the reason-giving that occurs in contexts such as request, orders, or similar practical scenarios. For example, if a police officer tells me to stop my car, this will give me a reason to do so

(and I probably ought to do it). Enoch uses the example of a request: if I ask you to read a draft of my paper, this gives you a reason to do so.

Robust reason-giving is clearly important in our practical lives, Enoch agrees. But how should we account for it? In terms of second-personal reason-giving, as Darwall claims? Does the power of requests or orders to generate reasons depend on the right of those making such demands and the accountability of the addressees? Enoch rejects that move and defends a reducibility claim instead: robust reason-giving is nothing but a particular instance of triggering reason-giving. While it appears as if there was something normatively special – perhaps second-personal – about robust reason-giving, properly understood this turns out to be misleading:

“when I ask you to read my paper, presumably there is this general reason (to do as I ask, within limits, in a certain context, etc.), one that I presumably succeed in triggering by making the request. It is in this way, then, that the suggested account of robust reason-giving is a particular instance (but an importantly unique one) of triggering reason-giving.” (Enoch 2011, 16)

Enoch’s argument for the reducibility claim invokes the distinction between a wide scope and a narrow scope reading of the conditional reason that I have to do as you ask me to (2011, 10).

Wide scope: You have a reason to (read the draft if I ask you to read it).

Narrow scope: If I ask you to read the draft, you have a reason to read it.

The wide scope reading shows how the reason-giving in question can be interpreted as a version of triggering reason-giving. The narrow scope reading, by contrast, does

not entail triggering reason-giving. If we only focus on the narrow scope reading, it appears as if robust reason-giving is distinct from triggering reason-giving. And if that were the correct account of robust reason-giving, it would be compatible, for example, with second-personal reason-giving.

Enoch argues, however, that while the narrow scope reading gives a plausible description of the practical phenomenon, it is only plausible in conjunction with a wide scope reading. The thought is that the wide scope reading explains the truth of the narrow scope reading and there is no alternative, better explanation of its truth. Since the wide scope reading characterizes triggering reason-giving, we have established the reducibility of robust reason-giving to triggering reason-giving. And if the reducibility claim is true, there is no need to resort to second-personal reasons to account for the normative phenomenon of robust reason-giving.

In addition, if the reducibility claim is true, then there is no further source of practical normativity – contrary to the claim that there are normatively distinctive, second-personal reasons. All normativity stems from the reasons that there are, i.e. from reasons that are not constructed and agent-relative in the fundamental sense that Darwall describes.

Responding to Normative Practical Uncertainties

Is the reducibility claim correct? And do the implications that Enoch draws from it for the prospects of Darwall's account of second-personal reason-giving hold? I want to grant to Enoch that some instances of robust reason-giving, including some that look

like instances that Darwall's account of second-personal reasons aims to cover, do indeed reduce to triggering reason-giving. Even Darwall's favored foot example may be a case in point. But granting such a weak interpretation of the reducibility claim – one that says that *some* instances of robust reason-giving reduce to triggering reason-giving – still leaves room for an account of second-personal reason-giving that is normatively significant. I reject the strong interpretation of Enoch's reducibility claim, according to which *all* instances of robust reason-giving reduce to a form of triggering reason-giving.

The account I shall offer diverges in an important respect from Darwall's original account of second-personal reasons, however. According to Darwall, second-personal reasons are fundamental for moral obligations. I propose to treat them as residual instead: second-personal reason-giving matters in contexts where third-personal reasons underdetermine how we ought to act. Specifically, I have in mind epistemic underdetermination. My argument, but not Darwall's, starts from the uncertainties that affect our practical deliberation, deliberation about how we ought to act. These uncertainties are particularly salient in interactional contexts and they can explain why second-personal reason-giving matters.

Practical uncertainties may be of a non-normative kind. We often do not have a completely accurate picture of the circumstances we are in and cannot anticipate with certainty all the consequences of our actions. For example, if I leave the house now and make my way to the station without delay, can I still catch the 2:30pm train to Oxford? There is a certain probability that I will catch the train and a certain

probability that I will not. This is not the kind of uncertainty I am concerned with here. The second kind of practical uncertainties is normative. It concerns the source of our normative beliefs and, as a result, the epistemic status of those beliefs: how can I know the difference between a true normative belief that correctly represents the normative reasons that apply and a belief that appears normative to me but that is simply the product of evolution or some other natural process? For example, is my inclination to offer help to someone supported by a correct belief that I have reason to help this person or is it merely a conditioned response that lacks warrant?

In his recent discussion of Derek Parfit's *On What Matters*, Darwall (2014, 91) argues that in light of this epistemological challenge as well as of our tendency to disagree in normative matters, doubt arises as to whether there are any normative facts at all. Parfit is aware of the problem, but defends the claim that practical normativity derives from normative facts. His argument hinges on a convergence claim. Parfit (2011) argues that our normative uncertainties and disagreements are merely superficial; underlying those disagreements is a substantive convergence among the three main normative theories – consequentialism, Kantian contractualism, and Scanlonian contractualism. This convergence is best explained, Parfit argues, as driven by the normative facts that the different theories attempt to track. If the convergence claim were true, we could identify correct beliefs about normative facts in the area of convergence among those theories and thus eliminate uncertainties. Darwall rejects Parfit's solution, however, and argues that the convergence claim is false (2014, 99ff).

If the convergence claim is false, and I agree with Darwall that it is, normative practical uncertainties (and the resulting disagreements) are back. But the correct response to normative practical uncertainties, I believe, is not to doubt the existence of normative facts altogether. The correct response, instead, is via a fuller development of practical justification under circumstances of uncertainty. When trying to establish how we ought to act, we often do not have knowledge of the reasons that apply. Vice versa, we may act on what we take to be reasons that apply even though our beliefs are false and hence we may not do what we have reason to do. Parfit is, of course, aware of this predicament. To capture the subjective element in practical deliberation, Parfit draws a distinction between “apparent” reasons and “real” reasons (2011, 35).

What is the normative status of such apparent reasons? Parfit answers this question by drawing a distinction between responding to reasons and being practically rational. Acting on apparent reasons is all that practical rationality requires, he claims:

“Our desires and acts are rational when they causally depend in the right way on beliefs whose truths would give us sufficient reasons to have these desires, and to act in these ways” (2011, 112).

Acting on apparent reasons may thus be permissible in the sense of being rational. But apparent reasons do not have normative force, only real reasons do (2011, 35).

Parfit further develops his view by distinguishing between different senses of ought (2011, 150f). He associates full-fledged normativity with a fact-relative sense of

ought – what we ought to do is determined by the real reasons that apply. The ought of practical rationality is belief-relative.¹ Finally, there is also an evidence-relative sense of ought, but that one plays no role in Parfit’s theory of practical normativity.

While I think that Parfit moves in the right direction here by recognizing how we often have to act without knowledge of the relevant normative facts, there is something puzzling about the idea that there are oughts that do not entail a normative reason. The problem with this way of understanding apparent reasons is that it is either incompatible with the claim that normativity is about reasons or with the claim that what we ought to do is normative (Kieseewetter 2012).

Part of this problem is Parfit’s way of characterizing practical rationality. This characterization is controversial, mainly because it focuses only on the belief-relative sense of ought and not the evidence-relative sense. I do not find it plausible. But since I am concerned here with the question of how we should understand normative reasons, this issue is tangential to the topic of my paper and so I want to bracket it.

Parfit’s understanding of normative reasons as giving rise to a fact-relative sense of ought is another part of the problem. While Parfit can allow for the possibility that the belief- and even the evidence-relative sense of ought influence our practical

¹ “[W]e ought rationally to act in some way when this act is what we ought practically to do in the belief-relative or normative-belief-relative sense” (2011: 163).

deliberation, they only do so at the level of apparent reasons. Normative reasons set a standard for the success of practical deliberation, but practical deliberation itself cannot give rise to normative reasons.

I do not think that this is right; this understanding of normative reasons rests on a truncated conception of our practical agency. In light of the practical uncertainties that surround practical deliberation, we need to introduce a further distinction, beyond the distinction between apparent reasons and real reasons. The distinction is between apparent reasons and constructed reasons. The thought is the following. There is an intuitive difference between a consideration that we mistakenly take to be a normative reason – because we hold a false belief – and a consideration that we take to have normative force in the absence of knowledge about the real reasons that apply to us. In Parfit's use of the term, both kinds of considerations fall under the category of apparent reasons and both may give rise to oughts, even though neither has the normative force of real reasons. But I do not see why we should elevate our simple mistakes in this way. Sure, we may sometimes be excused from making mistakes. But to be excused from making a mistake is not the same thing as acting as one ought to do, in any sense of ought. By contrast, it seems plausible to me that there are at least some responses to the uncertainty we face about which of our normative beliefs are true that can, as such, give rise to an ought. They cannot, of course, give rise to an ought in the fact-relative sense. But they nevertheless have independent normative force.

Based on this distinction we get three categories: (1) apparent reasons that are based on true beliefs about normative reasons – these are also real reasons, (2) apparent reasons that are given by mistaken normative beliefs and that do not have independent normative force, and (3) reasons that are constituted by a commendable response to the uncertainty we face about real reasons and that have independent normative force. I want to call the reasons of the third category constructed reasons. Constructed reasons are agent-relative.

A lot more would have to be said, quite generally, about the properties of a commendable response to normative practical uncertainties, but I cannot do this here. All I can do is discuss the rough idea with regard to the issue under debate, namely whether there are second-personal reasons or whether Enoch's reducibility claim about robust reason-giving is true.

Defending Second-Personal Reason-Giving

Enoch's alternative to Darwall's account hinges on the assumption that only what I have called, following Parfit, real reasons can have normative force in the sense of having the capacity to shape how we ought to act. As long as the relevant facts obtain, these reasons have normative force. Whether the relevant facts obtain depends on reason-triggering factors. To go back to Enoch's earlier example, once the pedestrian signals his intent to cross the road, the driver has reason to stop. The signal has triggered the relevant reason. Epistemic reason-giving, i.e. citing reasons for belief about the relevant facts, can help us form correct beliefs about how we

ought to act. If I am a passenger in your car, I might shout “stop – didn’t you see the pedestrian?” to make you aware of what you have – real – reason to do.

On Enoch’s account, robust reason-giving also needs to be explained in terms of real reasons. The interaction in those contexts in which robust reason-giving occurs is such that it triggers a latent real reason for action. For example, if my request that you read a draft of my paper truly triggers the conditional (read the draft if I ask you to read it), then my request succeeds in giving you a – real – reason to read it.

Importantly, on Enoch’s account, you have this reason independently of whether you are aware of the conditional or regard it as true. More generally, your attitudes towards the request are not part of the picture. Of course, my successful request does not imply that you now ought to read my draft – you may have much better things to do. But my successful request triggers a reason for action that changes your normative circumstances.

I want to grant Enoch that his account works for cases like the one just discussed. But I do not think that it works in all cases. First, consider what happens if the request is outlandish. Suppose I ask you to write my paper for me. I actually do not believe that you have a reason to write my paper for me and, I assume, neither do you. So we can agree that the conditional (write my paper if I ask you to write it) is false. If we are right, Enoch’s account implies that my request does not trigger a reason for action for you and your normative circumstances are thus unchanged.

Note that this implication is due to the fact that the conditional is false, not due to the fact that we agree that it is false.

Now consider the next case: what happens if we are uncertain about the truth of the relevant conditional? As it happens, I believe that this is not unusual; quite to the contrary, such uncertainties are ubiquitous and have a deep impact on our practical deliberation. How should we conduct our relationships with others? How can we establish whether our cherished goals deserve the attention we are inclined to give them? Should we obey an authority's directive that appears unjust or otherwise wrong? These are just a few examples of the kind of normative uncertainties that we often encounter. Suppose I am asking you to finish my paper for me because I find myself under unusual pressure, for example because I am currently suffering from some illness but my career hinges on that paper being completed. So the relevant conditional becomes something like (write my paper if unusual pressures force me to ask you to write it). This conditional (or a close cousin) might be true or it might not; I do not know. Suppose you do not know either.

On Enoch's account, all that we can say about a case like this is the following: if the conditional is true, then you have a reason to read my paper; if the conditional is not true, then you do not. If practical normative uncertainties are as ubiquitous as I believe they are, this response is somewhat unhelpful. I also think it is false. When we face uncertainties about the truth of the relevant conditionals, those normative uncertainties do not necessarily give rise to normative underdetermination. Qua practical agents, we are able to respond to those uncertainties and, in so doing, fill

the gap that is left by our lack of knowledge about the real reasons that apply to us. One way to put this is that responding to reasons is not the only function of practical deliberation. Deciding what to do when we do not know what reasons apply is another. With regard to the latter, the important thing to notice is this: some of these decisions will be better than others. But since we have already established that normative reasons underdetermine what we ought to do, we need a normative standard other than real reasons to make sense of this idea.

I want to propose that we think of second-personal reasons as setting such an alternative normative standard. Recall that Darwall's understanding of moral normativity in terms of second-personal reasons is in contrast to a fact-relative, third-personal understanding. What I am proposing is different: second-personal reason-giving does not replace third-personal reason-giving, not even in the moral domain. Instead, the two forms of reason-giving complement each other, given circumstances of normative practical uncertainty. In short, my basic idea is this: when we face uncertainties about the truth of a conditional, there is a second route to establishing its validity. This second route is second-personal and can give rise to constructed, non-fact-relative reasons for action.

As explained above, second-personal reason-giving presupposes mutually acknowledged relations of authority and accountability: your claim gives me a reason for action if we each accept that you have the authority to make that claim and that I am accountable to you in this regard. In his recent review of *On What Matters*, entitled "Agreement Matters", Darwall argues that Parfit is right to seek for

some sort of convergence or agreement about the normative reasons that apply. But he also argues that it is the second-personal standpoint, not the third-personal standpoint, that can account for the insight that an important dimension of normativity are “standards to which we justifiably hold ourselves and one another in common” (2014, 104).²

I want to take this view on board here and continue under the assumption that some form of agreement about what we can claim from each other can set a – second-personal – normative standard. Again, more would have to be said about what constitutes an agreement and what sort of disagreements undermine the possibility of creating second-personal reasons. This, too, will have to be done elsewhere. My modest aim here is to show how second-personal reason-giving may have independent normative force in circumstances of practical uncertainty. If it does, then Enoch’s strong reducibility claim is false.

How can second-personal reason-giving help us establish what we ought to do in circumstances of normative practical uncertainty? Before I can answer this question, I need to distinguish between different scenarios, relative to our epistemic circumstances. First, it might be that we agree that there is sufficient evidence for the truth of the conditional in question. We do not know that it is true, but we can agree that I have reason to write your paper. Second, we might agree on the opposite: in light of all the evidence we have, the conditional appears false and I do

² To be clear, Darwall’s claim – both in that review and in his 2006 book – is that the second-personal standpoint is necessary to account for the deontic dimension of morality. I am bracketing this stronger claim here.

not have a reason to write your paper. Again, we do not know that it is false, but we accept its falsity. Third, we might agree that the evidence is inconclusive and that the best response is that we should suspend belief. Fourth, we might disagree about what to believe about the conditional. We then need to respond to that disagreement and that response might again make suspension of belief the rational response. A fifth possible response is that we end up agreeing to disagree about the truth of the conditional.

In the last three cases, we cannot establish whether you have reason to write my paper based on our normative beliefs. But we can still establish this by agreeing on whether or not to uphold the conditional. If there is a reason for you to either write or not write the paper, it is because we can, independently of our beliefs about the truth of the conditional, agree that the conditional expresses a standard we can hold each other accountable to.

In all the cases just described, however, the answer to the question whether or not my request creates a reason for you depends on both our perspectives. Given that neither of us knows whether the conditional is true, the answer will depend, not just on our beliefs or our evidence, but also on our positive attempt to fill the normative gap created by our uncertainty, i.e. on the agreement we reach. If my request creates a reason for you, it is a constructed reason.

If we accept the possibility of constructed reasons, then we have left behind an account of reason-giving that is limited to epistemic and triggering reason-giving.

The weak reducibility claim may still be true, however: some robust reason-giving reduces to triggering reason-giving. But if robust reason-giving can involve constructed reasons, then the strong reducibility claim is false: not all robust reason-giving reduces to an instance of triggering reason-giving.

This account shares with Darwall's account the idea that the second-person standpoint is linked to the construction of reasons and hence to a distinctive source of practical normativity. But as already mentioned, my account of constructed reasons differs somewhat from Darwall's original account. The difference is that my uncertainty-driven account accepts the constraint of real reasons. The place of second-personal reason-giving is residual: it occurs as a response to the normative uncertainties we face. On Darwall's account, by contrast, there is something morally fundamental about second-personal reason-giving. My account is closer to Enoch's in this regard, as I accept the possibility of fact-relative oughts, even fact-relative moral oughts. Our interactions with others can determine our normative circumstances – second-personal reason-giving matters – but so can normative reasons all by themselves.

Acknowledgements

I am most grateful to David Enoch for extensive and very helpful written comments on an earlier version of this paper and for e-mail discussions of the main claims I am

trying to make. I have also benefitted greatly from Herlinde Pauer-Studer's and Kimberley Brownlee's suggestions.

References

Darwall, Stephen 2006: *The Second-person Standpoint*. Cambridge: Harvard University Press.

Darwall, Stephen 2014: "Agreement Matters." *Philosophical Review* 123(1), 79 – 105.

Enoch, David 2011: "Giving Practical Reasons." *Philosopher's Imprint* 11(4): 1 – 22.

Enoch, David 2014: "Authority and Reason-Giving." *Philosophy and Phenomenological Research* 89(2): 296 – 332.

Kiesewetter, Benjamin 2012: "A Dilemma for Parfit's Conception of Normativity." *Analysis* 72 (3), 466 – 474.

Korsgaard, Christine 1996: "The Reasons We Can Share." In Korsgaard (ed.) *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press, pp. 275 – 310.

Parfit, Derek. 2011: *On What Matters: Volume 1*. Oxford: Oxford University Press.

Pauer-Studer, Herlinde 2010: "The Moral Standpoint: First or Second Personal?" *European Journal of Philosophy* 18(2), 296 – 310.