



Agent-Relative Reasons and Normative Force

Jörg Löschke¹ 

Received: 16 May 2019 / Revised: 15 April 2020 / Accepted: 7 May 2020 /
Published online: 19 September 2020
© Springer Nature B.V. 2020

Abstract

The distinction between agent-relative reasons and agent-neutral reasons is philosophically important, but there is no consensus on how to understand the distinction exactly. In this paper, I discuss several interpretations of the distinction that can be found in the literature: the Motivational Interpretation, the Scope Interpretation, and the Goal Interpretation, and argue that none of these interpretations is entirely convincing. I propose a novel interpretation of the distinction, which I call the Normative Force Interpretation, according to which the distinction between agent-relative reasons and agent-neutral reasons concerns the normative force that these reasons can have for agents.

Keywords Agent-relative reasons · Agent-neutral reasons · Value theory · Restrictions · Relationships · Projects

The distinction between agent-relative reasons and agent-neutral reasons is widely recognized as philosophically important. Some authors even call it one of the most important findings in twentieth century ethical theory (Hurka 2003; Hammerton 2016). Among other things, it serves as a criterion to assess the plausibility of a moral theory: it is a common claim that consequentialism in its classical formulation cannot accommodate agent-relative reasons (McNaughton and Rawling 1991, 1993, 1995; Pettit 2000; Wallace 2009; Ridge 2017), and consequentialists themselves take this as an objection against their view, as evidenced by the fact that there have been numerous attempts to reconcile agent-relative reasons with a consequentialist framework. Very broadly speaking, the distinction matters because it expresses the idea that agents have a personal point of view from which some things matter to them more than others, and that moral theory should be able to accommodate our capability of taking a personal perspective on the world.

✉ Jörg Löschke
joerg.loeschke@lrz.uni-muenchen.de

¹ UZH Philosophisches Seminar, Zollikerstrasse 117, CH-8008 Zürich, Switzerland

However, despite its importance, there is no clear consensus on how to understand the distinction exactly. This comes to no surprise: the distinction is no part of our everyday language, and we cannot analyze the way regular people use the notions “agent-relative reasons” and “agent-neutral reasons” to determine how to understand the distinction correctly. It might even appear pointless to argue about how to understand the distinction. “Agent-relative reasons” and “agent-neutral reasons” are philosophical terms of art, and one might think that one can use them however one likes, as long as they are used consistently.

However, I think that this is a mistake. It is philosophically important and illuminating to reflect on the distinction and its best understanding. After all, the distinction has not been introduced for the sake of inventing philosophical terms of art, but to capture certain phenomena and their normative significance. And if we are interested in how moral theories can accommodate agent-relative reasons, then it certainly matters how to understand the distinction in the first place. In this paper, I will discuss several ways to interpret the distinction between agent-relative and agent-neutral reasons and will propose a novel interpretation that understands the distinction in terms of the normative force that reasons may or may not have. But first, some general remarks about the distinction are necessary.

1 General Remarks

A common way to characterize agent-relative reasons is to say that a reason is agent-relative if its full specification includes an ineliminable reference to the agent who has the reason, and that a reason is agent-neutral if it includes no such reference (McNaughton and Rawling 1998; Ridge 2005; Darwall 2006; Löschke 2014).¹ Following Nagel’s (1986) seminal work on the topic, most authors accept three kinds of agent-relative reasons, namely relationship-dependent reasons, project-dependent reasons, and deontological restrictions. For example, the fact that *my son* is drowning is an agent-relative relationship-dependent reason for me to save him; the fact that writing a chapter of a book on agent-relative reasons contributes to the completion of the book is a project-dependent agent-relative reason for me to write the chapter if it is *my project* to write a book on agent-relative reasons; and the fact that φ -ing would constitute a murder *by me* constitutes an agent-relative deontological restriction to not φ . There might be other kinds of agent-relative reasons,² but for the purposes of this paper, I will use this classical threefold distinction to discuss the best way of understanding the agent-relative/agent-neutral distinction.

An important aspect of agent-relative reasons is that agents often ought to act on their agent-relative reasons even if they could maximize (or minimize) actions of the very same type otherwise (Darwall 2006; Lippert-Rasmussen 2009; Nair 2014). Those who accept agent-relative reasons think that an agent ought to not kill an innocent, even if her killing the innocent would save three other innocents from being killed by others;

¹ McNaughton and Rawling (1998: 38) use the term “ultimate statement of the reason”, rather than “full specification”, but I take the idea to be the same, or at least very similar.

² Ridge (2017) emphasizes that egoistic reasons are also paradigm cases of agent-relative reasons. Dancy (1993) identifies seven kinds of agent-relative reasons in total.

that a father ought to take care of his children, even if neglecting his own children would cause three other fathers to take care of their children who would otherwise neglect them; and that agents ought to pursue their own projects instead of helping other to pursue their respective projects. Every plausible account of the agent-relative/agent-neutral distinction must be able to address this non-maximizing character of agent-relative reasons.

To see the intuitive difference between agent-relative reasons and agent-neutral reasons, it helps to reflect on the agent-neutral counterparts to the different kinds of agent-relative reasons. This will also be important for the subsequent discussion. The agent-neutral counterpart to my agent-relative reason to save my son is the fact that *a child* is drowning; the agent-neutral counterpart to my agent-relative reason to write the chapter for my book is the fact that *a person* has the project of writing a book on agent-relative reasons; and the agent-neutral counterpart to my agent-relative reason to not kill the innocent myself is the fact that *an innocent person is killed*. None of these reasons includes an essential reference to me as the agent, and this makes them agent-neutral.

An immediate worry with the distinction is that the definition of agent-relative reasons is merely formal: all it says is that in the case of agent-relative reasons, the complete reason statement includes a reference to the agent who has the reason. Why should such a formal definition matter? In fact, as McNaughton and Rawling (1991: 170–171) argue, merely pointing out that a reason includes an agential back-reference does not suffice to show that the reason is agent-relative, because agent-neutral reasons might also include a reference to the agent in their full specification. For example, a consequentialist might think that it is valuable in itself if parents look after their own children, so that the world as a whole is better if parents look after their own children. In such a case, a parent has a reason to look after her child, but not because this child is of particular importance to her, but because doing so promotes agent-neutral value.³ This suggests that, while drawing the agent-relative/agent-neutral distinction in terms of an agential reference within the reason statement is an important first step, it does not suffice to rely on linguistic considerations to understand the distinction.⁴ Instead, it is necessary to engage in substantive reasoning about how exactly we should understand the agential reference. Several interpretations are possible here, and I will discuss them in what follows.⁵

³ Strictly speaking, the agent-neutral reason would then be something like “the fact that a child is drowning (which happens to be mine)”.

⁴ This has also been pointed out by Krister Bykvist. See Bykvist (2018). However, Bykvist is more critical of the distinction than I am.

⁵ Another worry is that the different kinds of agent-relative reasons have nothing in common, and that there is not one way of drawing the distinction that applies to all relevant cases. After all, project- and relationship-dependent reasons differ from deontological restrictions in that the former speak in favor of performing certain acts, whereas the latter speak in favor of refraining from certain acts. Furthermore, reasons of the former kind sometimes permit an agent to put a special emphasis on her own *life*, and sometimes they require her to do so; reasons of the latter kind require the agent to put a special emphasis on her own *actions*. However, the assumption that there is unified interpretation of the distinction that applies to all kinds of agent-relative reasons seems warranted: in each case, the fact that something is mine – my project, my relationship, my action – is considered to be normatively relevant. The question is therefore in what way it matters that something is mine.

2 The Motivational Interpretation

Proponents of the Motivational Interpretation (MI) understand the reference to the agent in a motivational sense (Pettit 1987, 1988; Double 1999): the fact that something is mine motivates me to pursue a specific course of action. MI can be spelled out either in an externalist or in an internalist way. The externalist version holds that the reason-providing fact includes no agential reference, but the fact that I stand in a certain relation to that fact explains why I am motivated to do what I have reason to do. So, for example, the fact that *a child* is drowning gives me a normative reason to save the child, and the fact that the drowning child is *my son* motivates me to save him. The internalist version assumes that a reason statement can only be true if the reason can motivate the agent: an agent only has a reason to φ if there is a “sound deliberative route” (Williams 1995: 35) from the agent’s motivational set (her desires, commitments, loyalties, etc.) to φ . Hence, the fact that this particular child is *my son* motivates me to save him, and this motivational fact also constitutes a reason for me to save him. The externalist version thus holds that the agential reference motivates me to do what I already have reason to do, whereas the internalist version holds that the relative aspect not only motivates me, but also generates a reason for me: I have a reason to φ because the “my-ness” establishes the possibility that I am motivated to φ .

Neither version of MI is convincing. It is plausible to suppose that the reference to the agent is not only a motivating aspect, but also a normative consideration: it is part of the explanation why an agent ought to act in a specific way. A father who is not motivated to save his son is not thereby released from his duty to do so. Rather, he faces moral criticism precisely because he is not properly motivated. Furthermore, the agential reference is not necessary to explain motivation: the fact that *a child* is drowning can motivate me just as much as the fact that *my son* is drowning. Of course, my motivation to save my son might be stronger than my motivation to save a stranger (for example, I might be willing to risk my life to save my son, but not to save a stranger). But again, this cannot be the whole story about the agential reference. It does not explain why I *ought* to risk more to save my son, but merely describes the fact that I am in fact motivated to risk more. MI hence seems inadequate.

3 The Scope Interpretation

The internalist version of MI suggests a different interpretation of the agential reference, namely one that understands that reference as a reason-generating feature. On this view, the fact that *my son* is drowning generates a reason for me to save him. I will refer to this interpretation as the Scope Interpretation (SI), since it explains the difference between agent-relative reasons and agent-neutral in terms of the scope of the respective reasons. An agent-relative reason is a reason that applies only to one agent, whereas an agent-neutral reason is a reason that applies to everybody.

Many authors accept SI,⁶ and the interpretation does seem *prima facie* plausible. It explains the normativity of agent-relative reasons: a father who

⁶ See, for example, Parfit (1984); Mack (1998); Darwall (2002); Kolodny (2003); Huckfeldt (2007); Jeske (2008); Keller (2013). Schroeder (2007: 280) even calls this interpretation the “official definition” of the agent-relative/agent-neutral distinction. Lerm (2013) lists this interpretation as one possible interpretation without embracing it fully.

does not act on behalf of his son can be criticized because he does not act on a reason that he has due to his relationship with his son. And SI seems to yield the right results for other paradigm cases of agent-relative reasons. Take the case of project-related reasons. Intuitively speaking, I have a reason to finish a chapter for my book, but this is a reason that *I* have, given that writing this book is a personal project of mine. *You* do not seem to have such a reason, given that you do not plan to write such a book. In this sense, agent-relative reasons do seem to apply only to specific agents. By contrast, their agent-neutral counterparts apply to everyone. Suppose that a library collects every paper that has ever been written on agent-relative reasons, and that you have the choice between destroying the library or leaving it intact. Suppose that no other consideration is relevant for your choice except for the fact that *a person* has the personal project of writing a book on agent-relative reasons. In such a case, you certainly have a reason to leave the library intact. And this reason applies not only to you, but to everyone. SI thus seems to capture a relevant difference between agent-relative reasons and agent-neutral ones.

Nevertheless, SI is unsatisfactory. It holds that a reason is agent-relative if it applies only to the agent, and it is agent-neutral if it applies to everyone, but “this reason applies only to one agent” and “this reason applies to every agent” are not binary options. A reason can apply to several agents but fall short of applying to everybody. According to SI, such a reason is neither agent-relative nor agent-neutral, unless we think that every reason that does not apply to everybody is agent-relative. But we can imagine cases in which a reason applies to everybody except for one person. Does this reason cease to be agent-neutral simply because one person exists to whom the reason does not apply? This seems implausible.

The problem is especially troubling if we assume that the principle “ought implies can” applies not only to duties, but also to reasons, and that an agent only has a reason to φ if she can φ . And this seems to be a reasonable assumption: I arguably have no reason to fly to the moon, given that I am not capable of doing so. But if this is the case, then the distinction between agent-relative reasons and agent-neutral reasons collapses under SI. There are many actions that I cannot perform here and now, such as saving a drowning child in South Korea if I am currently in the US. If the principle “ought implies can” applies to reasons, then it follows that I have no reason to perform such acts. And this means that there are no agent-neutral reasons. For every reason, it is the case that the reason favors doing something that many agents cannot do due to contingent circumstances. Hence, no reason turns out to be agent-neutral, since there is no reason that everybody has.

Proponents of SI might respond that a reason is agent-neutral if it applies to every agent who is in a specific choice situation, whereas a reason is agent-relative if one agent would have this reason in a specific choice situation whereas other agents who are in the very same circumstances would not have this reason. They must then answer the question how to specify the relevant circumstances, and this is no easy task. Take the case of relationship-dependent reasons. Suppose that I stand by a lake and see my son drowning. How should we specify the relevant circumstances according to which my reason to save

my son is agent-relative? Maybe the correct description of the relevant circumstances is that I stand there and see my son drowning. But this view does not explain the agent-relative character of my reason, since every person in the circumstances thus described would have a reason to save my son from drowning: everybody who stands on that riverside and sees my son drowning has a reason to save him. Thus, if we specify the relevant circumstances in this way, the reason turns out to be agent-neutral, rather than agent-relative. Another way to specify the relevant circumstances is to include the fact that it is *my son* who is drowning in the description of the relevant circumstances. But this does not work, either. Again, every person in my circumstances would have the very same reason to save this particular child: *every* person who stands in the father-son-relationship to this particular child has the same reason to save him. Thus, my reason to save my son turns out to be agent-neutral again, because everybody in my circumstances has the same reason to save him. Finally, one might include in the description of the relevant circumstances not the fact that *my son* is drowning, but the fact that a child who stands to an agent in the father-child relation is drowning, whoever the father and the son are. But that does not solve the problem. Everybody who stands in this particular relation to the drowning child – and accordingly, everybody who is in my circumstances – has the same reason to save the drowning child.⁷

Thus, SI fails to achieve the right results. Either all reasons turn out to be agent-relative, or all reasons turn out to be agent-neutral. However, a possible variation of the Scope Interpretation might avoid this problem, namely what one might call the Goal Interpretation.

4 The Goal Interpretation

The Goal Interpretation (GI) interprets the agent-relative/agent-neutral distinction with reference to the goals or aims that agents have. This interpretation is inspired by the work of Derek Parfit. Parfit draws the agent-relative/agent-neutral distinction not only in terms of reasons, but also in terms of moral theories and understands a moral theory as agent-neutral if it gives all agents common aims, whereas a moral theory is agent-relative if it gives different agents different aims (Parfit 1984: 27). For example, consequentialism is an agent-neutral theory since it gives all agents the goal of making the world as good as possible, whereas deontology is an agent-relative theory since it gives each agent a different goal, such as the goal the *he* does not lie, kill, steal, etc. We can then say that a reason is agent-relative if it supports an action that aims

⁷ Of course, proponents of SI can respond that there is only one person who stands in the father-child-relation to this specific child, namely me, and that this suffices to show that nobody but me has this specific reason to save the child. I do not find this reply convincing. First, it might very well be that more than one person stands in the father-child-relation to this specific child (for example, I might be the social father, but somebody else is the biological father; or the biological father is deceased, but me and another person are both social fathers to this child). Second, the father-child-relation seems important because it is one variation of the parent-child-relation, but more than one person stands in the parent-child-relation to this specific child. I thank an anonymous referee for helping me to clarify this example.

at achieving a goal or realizing an end that moral theory gives specifically to that agent, and it is agent-neutral if it supports an action that aims at achieving a goal that moral theory gives to everyone.

GI might seem like a mere version of SI, since it also has the upshot that an agent-relative reason applies only to the agent, whereas an agent-neutral reason applies to everyone, but there is an important difference between GI and SI. According to GI, the my-ness by itself is not a reason-generating feature, strictly speaking. What generates the reason is the background theory that assigns moral goals to agents. GI understands the agential reference therefore more as an organizing feature than as a reason-generating feature. And GI also differs from the internalist version of MI because it does not rely on the goals and ends that agents actually have, but on the goals and ends that the moral background theory assigns them and that they ought to have. GI thus also explains the normativity of agent-relative reasons.

GI seems to sit well with the paradigm cases of agent-relative reasons. Project-dependent reasons arise if the background theory gives each agent the goal of pursuing her own project. My project of writing a book on agent-relative reasons gives me a reason to write a chapter for the book, and your project of running a marathon gives you a reason to go for a run on the weekend; these are reasons that we do not have in common, given that we have different projects and therefore different goals. Similarly, if moral theory assigns each agent the goal of taking care of her own child, then my goal is to ensure that my children get a good education, and your goal is to ensure that your children get a good education. Hence, I have an agent-relative reason to ensure that my children get a good education, and you have an agent-relative reason to ensure that your children get a good education. And GI can also make sense of restrictions: if moral theory assigns each agent the goal of not performing acts of certain kinds herself, then I have an agent-relative reason to not lie myself, even if my telling a lie would prevent three other agents from lying. After all, I have a different goal than these other agents. I do not have the goal to minimize the telling of lies (and neither do they), but the goal that *I* do not lie, whereas some other agent (call him Jack) has a different goal, namely the goal that *he* does not lie.

However, GI has more problems to deal with such cases than it initially appears. Talking in terms of goals invites a teleological view of practical reasons: what I have reason to do is to realize a goal, and this goal is plausibly understood as bringing about a certain state of affairs. GI thus claims that my agent-relative reason to not lie is a reason to bring about the state of affairs in which I do not lie. The problem is that deontological theories typically (and plausibly) hold that an agent not only has an agent-relative reason to not lie herself, but also an agent-neutral reason to prevent others from lying (for example, by reminding them that it would be wrong to tell a lie). And this means that Jack and I do have the same goal after all. Jack has a reason to not lie, and this is a reason to bring about the state of affairs in which he does not lie. But since I have a reason to prevent Jack from lying, I also have the goal to bring about the state of affairs in which Jack does not lie. States of affairs are not relativized or indexed to agents – there is no state of affairs_{Jack} or state of affairs_{me}. Thus, Jack and I do have the same goal,

namely, to bring it about that Jack does not lie. But if we both have the goal to bring about the state of affairs in which Jack does not lie, then it is difficult to see how GI can distinguish between agent-relative reasons and agent-neutral reasons.⁸

This problem generalizes to project- and relationship-dependent reasons. The plausible Kantian idea that agents have moral reasons to make the ends of others their own ends suggests that you and I share the same goals regarding our project-dependent reasons. We both have the goal to bring about the state of affairs in which I write a book on agent-relative reasons as well as the state of affairs in which you run a marathon. And we both have the goal to realize the state of affairs in which your children get good education as well as the goal to realize the state of affairs in which my children get a good education.

This is not to deny that agent-relativity allows for conflict. I ought to take care of my children's education, even if that means that you cannot take care of your children's education, and vice versa; I ought to not tell a lie, even if I could prevent you from lying, and vice versa. Such cases do not show that we have different goals but illustrate that the goals that we both share have a different impact on our normative situation. My reason to realize the state of affairs in which I do not lie has more impact on what I am required to do than my reason to realize the state of affairs in which you do not lie, and vice versa. And the same goes for the other cases of agent-relative reasons.⁹

This suggests yet another interpretation of the agent-relative/agent-neutral distinction which I will call the Normative Force Interpretation (NFI). This is what I will turn to now.

⁸ An anonymous referee has pointed out to me that agents may aim at other things than bringing about states of affairs, and that this may hurt my argument. While the question whether goals are state of affairs or something else this is a complex question that I cannot discuss in the detail it deserves here, I would like to say a couple of things in response.

First, I think that it is intuitive plausible to understand goals as states of affairs that agents aim to bring about rather than actions that they aim to perform. Suppose that my goal is that my children have financial security after my death. It does not seem obvious that this is the same goal as my goal to do what I can do to ensure that my children have financial security after my death. Or suppose that my goal is to be a good husband. This also does not seem to be the same as my goal to do everything that a good husband does.

A second consideration to think of goals as states of affairs comes from action theory. So-called teleological accounts of actions or of reasons for action understand actions as "means by which we affect how the world goes" (Portmore 2011: 56; for a defense of a teleological account of action, see also Sehon 1997). And, as Portmore points out, this means that acting consists in actualizing possible worlds. And actualizing possible worlds is plausibly understood as bringing about states of affairs. Now, of course, the teleological conception of either actions or of reasons for action is controversial, but it provides independent support for the idea that goals are states of affairs.

Third, when an agent pursues a goal, she considers the thing that she aims at to be good in some form. Thus, we can expect the question of what goals are – states of affairs, or actions, or perhaps something else – to mirror the debate on the fundamental bearers of value. And while there are many different accounts about the correct understanding of fundamental bearers of value (for an overview of the possible positions here, see Olson 2004: 32), states of affairs seem to be a plausible candidate for a fundamental bearer of value; and as Wedgwood (2009: 327) has pointed out, it is possible to translate every claim about value into a claim about a valuable state of affairs. If this is correct, then we should be able to translate every claim about an agent's goal into a claim about a state of affairs that she aims to bring about.

⁹ An anonymous referee has questioned whether it really makes sense in such a case to say that we both share the same goal. Priority between more general aims seems to affect what each of us should aim at when making a particular choice. Nevertheless, I think it still makes sense to say that agents share the same goals, even if they prioritize the shared goals in different ways. In other contexts, it seems plausible to say that you still have a specific goal (say, start a family), even though some other goal of yours has priority at that moment (say, advancing your career). Similarly, we might say that in situations of conflict, agents still have the same goals even though they prioritize them differently. But proponents of GI do not claim that in cases of agent-relative reasons we have different priorities among our goals; they say that we have different goals.

5 The Normative Force Interpretation

According to NFI, the fact that *my* daughter is drowning, or that writing a chapter for the book helps to facilitate the success of *my* project should not be understood in an ontological or metaphysical way, in the sense that it somehow generates a reason that I would not have otherwise. Instead, it should be understood as a feature that affects the normative force or the strength of a reason that I already have. Accordingly, to say that the fact that my daughter is drowning gives me an agent-relative reason to save her includes two things: there is a fact that gives me a reason (such as the fact that a person is drowning), and the normative force that this reason has for me cannot be understood without taking the relationship between me and my daughter into account.¹⁰

NFI is hardly spelled out in the literature, but some authors seem to suggest it. For example, Theo van Willigenburg writes that “[a]gent-relative reasons, in distinction from agent-neutral ones, are reasons the authority of which cannot be understood without an essential reference to the relationships or the projects of a particular agent” (van Willigenburg 2005: 45 f.), and R. Jay Wallace writes that agent-relativity is the phenomenon “whereby considerations provide one agent with reasons to do something, without having the same kind of normative significance for other agents who are differently situated” (Wallace 2010: 519). Saying that the reference to the agent explains the authority of a specific reason is not the same as saying that the reference to the agent explains the existence of the reason: a reason that is generated by something else can also have different authority over different agents. And saying that considerations provide one agent with reasons to do something without having the same kind of normative significance for others can mean either that these considerations do not give other agents reasons, or it can mean that the considerations give other agents reasons, but with a different normative force.¹¹

I do not claim that van Willigenburg or Wallace embrace NFI, but their remarks do suggest an interpretation of the agent-relative/agent-neutral distinction along these lines. On NFI, a reason is agent-relative if it is not possible to understand the normative force of the reason for the agent without reference to the agent, and a reason is agent-neutral if its normative force does not depend on the identity of the agent. Thus, a reason is agent-relative if its normative force can differ from agent to agent, and it is agent-neutral if its normative force remains constant among agents.

An immediate worry with NFI is that, while it is plausible in some cases, it yields implausible results in other cases. According to NFI, every agent-relative reason that A has is also a reason that B has, albeit with less normative force. This might be plausible with regard to certain relationship-dependent reasons: if my daughter is drowning, I have an agent-relative reason to save her. That reason does not exist because of our relationship but because of her predicament, and strangers therefore also have a reason

¹⁰ To repeat, it does not suffice to say that the reason is agent-relative simply because it is my daughter who is drowning, for reasons that include an agential back-reference can also be agent-neutral reasons. The reference to the agent needs to do normative work for the reason to count as agent-relative and altering the force of my reason to save this child is such normative work. I thank an anonymous referee for pressing me on this.

¹¹ See also Bykvist (2018) who is critical about the agent-relative/agent-neutral distinction and thinks that we should rather focus on how the weighing of reasons should be done according to different moral theories. I am obviously sympathetic to this view, but I think that NFI shows that the agent-relative/agent-neutral distinction matters quite a bit when we talk about weighing reasons according to different moral theories.

to save her (if they can do so). But of course, the reason to save my daughter figures much stronger in the explanation of what I ought to do or what I am required to do than in the case of strangers. It stems from her predicament, but it has more normative force for me than it has for strangers, due to my special relationship with her. But in other cases, NFI might seem less plausible. Take the case of restrictions. I have a reason to keep a promise, and this is a reason for *me* to *do* something. But how can you have that reason? You cannot perform my actions for me; hence, it might seem odd to say that you and I have the same reason that merely differs in normative force. You can help me keeping my promise, but your reason to help me keeping my promise is not the same reason as my reason to keep my promise.

This is a fair objection, but it does not refute NFI, as the discussion of GI above has shown. My reason to keep my promise and your reason to help me keep my promise can be both understood as reasons to bring about the state of affairs in which I keep my promise.¹² And on this understanding, it is much more plausible to say that you and I do have the same reason, with the difference being that my reason has more normative force for me than your reason has for you, even if you cannot perform my action.

Project-dependent reasons might also seem to pose a problem for NFI. If your project is to run a marathon, then you have project-dependent reasons to do certain things, such as going for a run on Sunday. To say that this is an agent-relative reason is to say that not only do you have a reason to promote the state of affairs in which you successfully finish the marathon, but every other agent has a reason to help you doing so. And this might appear odd. Some readers might think that you have a reason to train for the marathon, but that your project gives others no reasons whatsoever: an agent's personal projects are her personal matter, and your project-dependent reasons therefore do not apply to other agents.

However, I do not think that this objection to NFI is very convincing. Recall the aforementioned Kantian idea of making the ends of others our own ends. If you think it is plausible that this is a plausible idea – and I do not see any reason to think that morality does not include the idea that we all have reason to help each other in achieving our individual ends – then the idea that everybody has reason to help me finish the marathon is less odd than it might initially appear. Of course, nobody has a moral duty to help others achieve their projects, but agents certainly have reason to do so. Suppose that a stranger's project is to run a marathon, but she lacks the money to buy new running shoes that she desperately needs, and I somehow find out about this. It seems plausible to suppose that I have an altruistic reason to buy these shoes for her (if I can do so). Of course, I am probably not required to buy the shoes for her – my reason to buy the shoes is easily outweighed by other reasons that I have. But that does not show that the reason does not exist – it only shows that the reason has little normative force for me, much less force than her reason to do everything she can to finish the marathon – and this is just what NFI claims.

NFI has several advantages. It avoids problems that the other interpretations of the agent-relative/agent-neutral distinction face; it fulfills the desiderata that a convincing interpretation of the distinction must fulfill; and it also has further theoretical advantages. I will address these points in the next section.

¹² See my remarks in FN 7.

6 Some Advantages of NFI

The first advantage of NFI is that it avoids the problems that the other interpretations face. In contrast to MI, NFI can explain the normativity of agent-relative reasons, since it presupposes that agent-relative reasons are normative reasons, not merely motivating reasons. SI understands the agential reference as a reason-generating feature; as a result, it cannot explain why associates and non-associates often have the same practical reasons to act on someone's behalf. NFI avoids this problem. It presupposes that agent-relative reasons are grounded in agent-neutral considerations and can therefore easily explain why an agent has a reason to φ both on behalf of an associate and on behalf of a stranger. This also means that NFI does not face SI's problems to specify the relevant circumstances under which agent-relative reasons arise. GI, in its most plausible form, also operates with the idea of differences in the normative force that reasons can have, and hence collapsed into NFI.

At the same time, NFI captures the phenomena that motivate the distinction in the first place just as much as the other interpretations do. NFI is in line with the idea that agent-relative reasons stem from the first-person perspective of agents and that moral theory must make room for the importance of what an agent considers to be important and valuable in life. In addition, it explains the non-maximizing logic of agent-relative reasons: if your reason to not kill an innocent yourself has special normative force for you, then this explains why that reason outweighs your reason to kill an innocent in order to save three.

NFI has the further advantage that it can explain better than other interpretations why restrictions are not absolute and why, at some point, it might be justified to kill a person to save a sufficiently high number of persons. At some point, your reason to not kill a person yourself might be outweighed by your reason to save many others, its special normative force notwithstanding. This seems to be the right result, and it seems more difficult to explain it on SI or GI: these other interpretations presuppose that agent-relative reasons and agent-neutral reasons differ in kind, and not merely in normative force. Their explanation of why I ought to not kill an innocent to save even more innocents from being killed thus seems to be that agent-relative reasons cannot be weighed against agent-neutral reasons, or that agent-relative reasons override agent-neutral reasons. But if that is the case, how can it be ever be the case that agent-neutral reasons outweigh agent-relative reasons? Proponents of SI or GI might be able to come up with a plausible explanation, but *prima facie*, NFI gives a more natural explanation.

Finally, NFI can answer to objections that have been raised not only against SI or GI, but also objections that have been raised against the agent-relative/agent-neutral distinction in general. Again, SI or GI might also be able to answer these objections, but the fact that NFI *can* answer to these objections, together with the fact that NFI can make sense of the phenomena that motivate the distinction in the first place and the fact that it does not face the same problems that other interpretations face suggests that NFI is the best account of the agent-relative/agent-neutral distinction.

A first objection against the agent-relative/agent-neutral distinction comes from Korsgaard (1993). Korsgaard argues that the distinction cannot make sense of deontological restrictions. Restrictions must be grounded in reasons that are shareable in the sense that not only does the agent have a reason to not perform a certain act, but the potential victim must also be able to appeal that reason and demand compliance with it,

and Korsgaard thinks that this is not possible on an account that grounds restrictions in agent-relative reasons: “If the deontological reasons were agent-relative, merely my property, my victim would not have the right to demand that I act on it” (Korsgaard 1993: 27).

However, Korsgaard is wrong in assuming that agent-relative reasons must somehow be the personal property of agents, and NFI explains why. Agent-relative reasons could only be the private property of agents if the my-ness somehow generates the reason. Since NFI rejects this assumption, it also rebuts Korsgaard’s attack on the agent-relative/agent-neutral distinction. Deontological restrictions are agent-relative reasons, but they are grounded in goals that are shared by everyone. My-ness does not generate reasons but increases the normative force of reasons that exist anyway, and this explains why agent-relative reasons are necessarily shared reasons: the very same feature that gives me a reason gives every other agent that reason as well. Hence, there is no problem to explain why potential victims of my actions can demand my compliance with my deontological reasons.

Rønnow-Rasmussen (2009) has raised another objection against the agent-relative/agent-neutral distinction. He argues that all reasons are agent-relative on purely formal grounds, because a reason is always a reason for someone. States of affairs can be good without reference to any agent, but reasons do not work this way: to say that there is a reason to φ is to say that someone has a reason to φ . And this means that the specification of *every* reason must include a reference to the agent who has the reason. Hence, every reason is agent-relative, and the distinction turns out to be pointless.

Against this objection, NFI shows that the distinction is not pointless. The fact that every reason is a reason for someone is neither here nor there, because it does not imply that every reason has the same normative force for everyone. Rønnow-Rasmussen’s objection shows that merely identifying an essential reference to an agent within the reason statement is not enough to understand the agent-relative/agent-neutral distinction; we also need to understand what substantial role the agential reference plays. NFI provides a substantial account of the agential reference, and this means that not every reason turns out to be agent-relative on purely formal grounds.

Another objection against the agent-relative/agent-neutral distinction comes from Thomas (2003). Thomas emphasizes the way how agent-relative reasons relate to the agent’s point of view: agent-neutral reasons arise from an impersonal perspective and correspond to agent-neutral value, and agent-relative reasons arise once an agent switches from that impersonal perspective to a personal point of view from which the agent not only regards things as impersonally valuable, but also regards things as constituents of her good life and sees what she cares about. If agent-neutral reasons correspond to agent-neutral value, then by analogy agent-relative reasons seem to correspond to agent-relative value. And this might lead to an implausible result, because it implies that we add value to the world simply by switching from an impersonal to a personal point of view: by making the switch, we introduce agent-relative value to the world that wasn’t there before. This is certainly metaphysically dubious.

By now it should be clear that NFI does not imply such metaphysical extravaganzas. NFI does not imply that the switch from the impersonal perspective to the personal perspective generates new value. It merely relies on the observation that certain values are more salient to an agent than others. NFI does not claim that agent-relative reasons

correspond to agent-relative value: it claims that there is only one kind of value, namely agent-neutral value, and that this kind of value generates reasons of two sorts, namely agent-relative as well as agent-neutral reasons. If all agents have the same reasons, and if the difference between agent-relative reasons and agent-neutral reasons is a difference in the normative force of these reasons, then it is quite natural to suppose that these reasons have the same (agent-neutral) source. Hence, there is no metaphysically suspicious generation of value by switching from the impersonal to the personal perspective.

7 Concluding Remarks

The distinction between agent-relative reasons and agent-neutral reasons is an important distinction that expresses the idea that as valuing beings, some things are more important to us than others. Any plausible moral theory should make room for agent-relative reasons. But it matters how the distinction is to be understood, because this has an impact on the practical upshot of our first-personal perspective. I have argued that the distinction should be understood in terms of normative force. While many things provide us with a multitude of reasons for action, some of these reasons have more normative force on us than others.

Of course, more needs to be said here. For example, it is necessary to explain why my-ness alters the normative force of reasons, and how exactly this is to be understood. I cannot provide an explanation here, but I have discussed it elsewhere (Löschke 2020). But providing a plausible interpretation of the agent-relative/agent-neutral distinction is by itself theoretically important, regardless of the explanation of how agent-relative reasons are possible.

Acknowledgements This research has been funded by the Swiss National Science Foundation as part of the research project on “Value-Based Non-Consequentialism” (Grant Number PP00P1_176703 / 1). For helpful comments and feedback, I thank Monika Betzler, Julian Nida-Rümelin and Christoph Knill, as well as a very helpful anonymous referee.

References

- Bykvist, K. (2018). Agent-relative and agent-neutral reasons. In D. Star (Ed.), *The Oxford handbook of reasons and normativity* (pp. 821–838). Oxford: Oxford University Press.
- Dancy, J. (1993). *Moral reasons*. Oxford: Blackwell.
- Darwall, S. (2002). *Welfare and rational care*. Princeton: Princeton University Press.
- Darwall, S. (2006). *The second-person standpoint. Morality, respect, and accountability*. Cambridge: Harvard University Press.
- Double, R. (1999). Impartiality, and what we can ask of persons. *American Philosophical Quarterly*, 36, 149–158.
- Hammerton, M. (2016). Patient-relativity in morality. *Ethics*, 127, 6–26.
- Huckfeldt, V. (2007). Categorical and agent-neutral reasons in Kantian justification of morality. *Philosophia*, 35, 23–41.
- Hurka, T. (2003). Moore in the middle. *Ethics*, 113, 599–628.
- Jeske, D. (2008). *Rationality and moral theory. How intimacy generates reasons*. New York: Routledge.
- Keller, S. (2013). *Partiality*. Princeton: Princeton University Press.

- Kolodny, N. (2003). Love as valuing a relationship. *Philosophical Review*, 112, 135–189.
- Korsgaard, C. (1993). The reasons we can share: An attack on the distinction between agent-relative and agent-neutral values. *Social Philosophy and Policy*, 10, 24–51.
- Lerm, J. (2013). The agent-relative/agent-neutral distinction: My two sense(s). *South African Journal of Philosophy*, 32, 137–148.
- Lippert-Rasmussen, K. (2009). Kamm on inviolability and agent-relative restrictions. *Res Publica*, 15, 165–178.
- Löschke, J. (2014). Partiality, Agent-Relative Reasons, and the Individuals View. *Social Theory and Practice*, 40(4), 673–682.
- Löschke, J. (2020). Agent-Relative Reasons as Second-Order Value Responses. *Canadian Journal of Philosophy*, 50(4), 477–491.
- Mack, E. (1998). Deontic restrictions are not agent-relative restrictions. *Social Philosophy & Policy*, 15, 61–83.
- McNaughton, D., & Rawling, P. (1991). Agent-relativity and the doing-happening distinction. *Philosophical Studies*, 63, 167–185.
- McNaughton, D., & Rawling, P. (1993). Deontology and agency. *The Monist*, 76, 81–100.
- McNaughton, D., & Rawling, P. (1995). Value and agent-relative reasons. *Utilitas*, 7, 31–47.
- McNaughton, D., & Rawling, P. (1998). On defending deontology. *Ratio*, 11, 37–54.
- Nagel, T. (1986). *The view from nowhere*. New York: Oxford University Press.
- Nair, S. (2014). A fault line in ethical theory. Consequentialism, deontic constraints, and the Prisoner's dilemma. *Philosophical Perspectives*, 28, 173–200.
- Olson, J. (2004). Intrinsicism and Conditionalism about final value. *Ethical Theory and Moral Practice*, 7, 31–52.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Clarendon Press.
- Pettit, P. (1987). Universalizability without utilitarianism. *Mind*, 96, 74–82.
- Pettit, P. (1988). The paradox of loyalty. *American Philosophical Quarterly*, 25, 163–171.
- Pettit, P. (2000). Non-Consequentialism and Universalizability. *The Philosophical Quarterly*, 50(199), 175–190.
- Portmore, D. W. (2011). *Commonsense Consequentialism*. Wherein Morality Meets Rationality. New York: Oxford University Press.
- Ridge, M. (2005). Universalizability for collective rational agents: A critique of agent-relativism. *Philosophy and Phenomenological Research*, LXX, 34–66.
- Ridge, Michael (2017): Reasons for action: Agent-neutral vs. agent-relative. In: Edward N Zalta (ed.): The Stanford Encyclopedia of Philosophy (Fall 2017 Edition). <https://plato.stanford.edu/archives/fall2017/entries/reasons-agent/>. Accessed 06 May 2020.
- Rønnow-Rasmussen, T. (2009). Normative reasons and the agent-neutral/agent-relative dichotomy. *Philosophia*, 37, 227–243.
- Schroeder, M. (2007). Reasons and agent-neutrality. *Philosophical Studies*, 135, 279–306.
- Sehon, S. (1997). Deviant causal chains and the irreducibility of teleological explanation. *Pacific Philosophical Quarterly*, 78, 195–213.
- Thomas, A. (2003). Nagel's paradox of equality and partiality. *Res Publica*, 9, 257–284.
- van Willigenburg, T. (2005). Reason and love: A non-reductive analysis of the normativity of agent-relative reasons. *Ethical Theory and Moral Practice*, 8, 45–62.
- Wallace, R.J. (2009). The publicity of reasons. *Philosophical Perspectives*, 23, 471–497.
- Wallace, R. J. (2010). Reasons, values and agent-relativity. *Dialectica*, 64, 503–528.
- Wedgwood, R. (2009). Intrinsic values and reasons. *Philosophical Issues*, 19, 321–342.
- Williams, B. (1995). *Making sense of humanity*. Cambridge: Cambridge University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.