Canadian Journal of Philosophy (2020), 50: 4, 477–491
doi:10.1017/can.2020.3

ARTICLE

Agent-Relative Reasons as Second-Order Value Responses

Jörg Löschke

University of Zurich, Department of Philosophy, Zurich, Switzerland
Email: joerg.loeschke@philos.uzh.ch

Abstract
Agent-relative reasons are an important feature of any nonconsequentialist moral theory. Many authors think that they cannot be accommodated within a value-first theory that understands all value as agent-neutral. In this paper, I offer a novel explanation of agent-relative reasons that accommodates them fully within an agent-neutral value-first view. I argue that agent-relative reasons are to be understood in terms of second-order value responses: when an agent acts on an agent-relative reason, she responds appropriately to the agent-neutral value of her own appropriate response to some agent-neutral value. This view helps reconcile important elements of deontology and consequentialism.

Keywords: Agent-relative reasons; agent-neutral value; restrictions; relationships; deontology; consequentialism; appropriate value responses

1. Introduction

In recent years, there has been extensive discussion about the relation between reasons and values. Many authors endorse the so-called Buck-Passing Account of Value (BPA) (Scanlon 1998). According to BPA, facts about reasons explain facts about goodness: an object's property of being good is the second-order property of having other properties that give agents reasons to respond with pro-attitudes to the object (Schroeder 2009; Rønnow-Rasmussen 2011; Rowland 2013). BPA has several advantages: it is metaphysically parsimonious (Suikkanen 2009), it demystifies value (Rabinowicz and Rønnow-Rasmussen 2004; Lang 2008), and it explains the normativity of value (Dancy 2000; Heuer 2006; McHugh and Way 2016). However, it also faces serious objections. A well-known objection is the wrong kind of reason problem (Rabinowicz and Rønnow-Rasmussen 2004, 2006; Danielsson and Olson 2007; Lang 2008; Rowland 2013; Kauppinen 2014a): agents can have reasons for pro-attitudes towards objects that are not good. Furthermore, different values can involve reasons for the same kinds of response, and proponents of BPA have difficulties explaining how we can discriminate between such values (Crisp 2005; Väyrynen 2006). Another problem for BPA is that the deontic realm is not coextensive with the evaluative realm because a reason is always a reason for someone, but objects can be good independently of any valuing agent (Dancy 2000; Bykvist 2009; Hurka 2014). Finally, it has been objected that reasons cannot be the normative primitive because they can be evaluated themselves (Gregory 2014, 2016).

These apparent problems have motivated new research on the prospects of value-based theories of reasons (Orsi 2013; Maguire 2016). According to such a theory, facts about value explain facts about reasons. Value-based theories of practical reasons are well suited to avoid the objections against BPA. However, they also face problems. This paper is part of the broader research project of exploring the prospects of value-based theories of practical reasons. Of course, it is not possible to defend such a theory in detail in one paper. What I aim to do is defend a value-first view against one objection in particular—namely the problem of agent-relative reasons.

The problem is this. Once it is assumed that a value-based theory of practical reasons is plausible, it is also plausible to suppose that all value is agent-neutral. Value is agent-neutral if a relation of the valuable entity to a specific agent is no fundamental good-making property, and this is an intuitively plausible view. For example, what makes my pleasure or the well-being of my children good is not the fact that it is *my* pleasure or the well-being of *my* children—that would be quite an eccentric view. Rather, pleasure or the well-being of persons is good, and this goodness can occur in my life or the life of my children. Agent-neutrally good entities can be good for agents, and the existence of an agent-neutrally good entity can depend on a specific agent (my pleasure is plausibly good for me, and it only exists because I exist), but the goodness of an agent-neutrally valuable entity does not depend on the identity of the agent.[1] Now, if facts about value explain facts about reasons, and if the goodness of an entity does not depend on the identity of any particular agent, then it seems natural to suppose that value generates the same reasons for all agents. But commonsense morality acknowledges so-called agent-relative reasons—reasons that differ from agent to agent. Suppose that your daughter and my daughter are both drowning, and that there is only one life preserver. Both your daughter drowning and my daughter drowning are bad states of affairs and, on a value-based theory of reasons, this means that both states of affairs give us both reason to throw the life preserver into the water. They might even give us both reason to throw it to any of these girls. But commonsense morality clearly suggests that I have a special reason to throw the life preserver to my daughter, and you have a special reason to throw it to your daughter. But if value is understood as agent-neutral value, if the goodness of your daughter's survival has nothing to do with *you* in particular, and if the goodness of my daughter's survival has nothing to do with *me* in particular, then it is difficult to see why these states of affairs give us different reasons. The question, then, is how such agent-relative reasons can be accommodated within a value-based theory of reasons that understands value as agent-neutral.

Many authors doubt that this is possible. Some conclude that agent-relative reasons are not grounded in value and that they therefore speak against value-based theories of reasons (McNaughton and Rawling 1992; Jeske 2008; Wallace 2010; Chappell 2014). Others propose abandoning an agent-neutral understanding of value and adopting an agent-relative understanding of value instead (Smith 2003, 2009; Louise 2004; Portmore 2011).[2]

Against these authors, I will offer a novel account of how to accommodate agent-relative reasons within an agent-neutral value-first framework. I will explain agent-relative reasons in terms of second-order value responses: when an agent acts on her agent-relative reasons, she responds appropriately to her own appropriate value response. The gist of this paper is expository: I will explain the view, rather than discuss alternative views. However, since it offers an answer to an

[1]In recent years, there have been several accounts of "value" that make the idea of "agent-neutral value" appear problematic, especially context-sensitive accounts of goodness. Examples include Stephen Finlay's (2014) account that understands goodness claims as relative to ends, and Michael Ridge's (2014) account that understands goodness claims as relative to standards. According to such views, value is always relative (to ends, or standards, or whatever), and since the relevant standards or ends might include pronominal back-reference, the assumption that all value is agent-neutral value might appear problematic.

I cannot discuss here these context-sensitive accounts of goodness in the detail they deserve. However, some points deserve mention. Barry Maguire (2016, 234n2) points out that Finlay's view allows for quite idiosyncratic ends, and that his view therefore counterintuitively allows for bad or worthless ends to generate reasons. And Rowland (2016) puts forward some arguments against Ridge in favor of the intuitively plausible view that entities can be good *simpliciter* rather than context-sensitively good. Following these arguments, I think the idea of agent-neutral value has the important advantage of being intuitively compelling. It is also in line with the phenomenology of value experience: when we experience something as good, we typically experience it as not entirely up to us whether the thing is good, and we often experience it as good *simpliciter,* rather than as good-relative-to-some-*X*. It would require more work to defend the agent-neutral account of value in detail. In any case, its intuitive persuasiveness makes it a view that is at least worth exploring. I thank an anonymous referee for pressing me on this point.

[2]For some problems with the notion of agent-relative value, see Schroeder (2007a). These problems also support an agent-neutral understanding of value.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

important objection against an agent-neutral value-based theory of practical reasons, it also speaks in favor of such a view, at least indirectly.

The structure of this paper is as follows. First, I lay out what a plausible version of a value-based theory of reasons might look like. I then explain the agent-relative/agent-neutral distinction in some detail. Subsequently, I present the view of explaining agent-relative reasons in terms of second-order value responses. I conclude by briefly mentioning some advantages of the view.

2. Reasons as appropriate value responses

Value-based theories of practical reasons are usually associated with consequentialism.[3] Consequentialists understand practical reasons as reasons to promote valuable states of affairs. Hence, consequentialists interpret the value-based theory of reasons in the following way:

> *Consequentialist Reasons* (CR): An agent has a reason to φ iff (and because) her φ-ing would promote some valuable state of affairs *S*.[4]

CR has two important features. First, it is a teleological conception of practical reasons (Portmore 2011; Nye, Plunkett, and Ku 2015). It understands practical reasons as goal oriented in the sense that they are reasons to bring about possible worlds. Second, according to CR, states of affairs are the bearers of value. The value of *S* provides reasons for action, not some object that constitutes or is part of *S*.

Neither of these features is a necessary commitment for a value-based theory of reasons. It is also possible to interpret practical reasons more broadly as reasons for appropriate value responses:

> *Reasons as Appropriate Value Responses* (AVR): An agent has a reason to φ iff (and because) her φ-ing would constitute an appropriate response to some agent-neutral value *V*.

AVR can be interpreted in a way that is identical to CR: if *V* stands for the value of a state of affairs *S* and if the appropriate response to *V* is to promote *S*, then AVR and CR amount to the same view. But AVR is also compatible with a pluralist account of value bearers, as well as with the view that other kinds of value responses than promotion can be appropriate. For example, AVR is compatible with the Kantian view that persons are bearers of value, and that the appropriate response to the value of a person is to respect her. Since AVR is broader and compatible with many different views, I will rely on it in what follows.[5]

To see in what sense agent-relative reasons pose a problem for AVR, let me now elaborate on the distinction between agent-relative reasons and agent-neutral reasons in some detail.

3. The agent-relative/agent-neutral distinction

The distinction between agent-relative reasons and agent-neutral reasons is one of the most important findings in twentieth-century ethical theory (Hurka 2003; Hammerton 2016). Among other things, it is a criterion for the plausibility of a moral theory.[6] A theory that cannot accommodate agent-relative reasons is deficient for that very reason.[7]

[3]Some authors identify consequentialism with the attempt to explain reasons in terms of goodness. See, for example, Nair (2014).

[4]Many consequentialists also hold a specific view of moral requirements: agents are required to perform those actions that will lead to the best outcome overall.

[5]I cannot offer a full defense of AVR here. My point is simply that AVR does not settle a number of other controversial matters. This makes it a good working hypothesis for proponents of a value-based theory of reasons.

[6]See Ridge (2011) for further ways in which the distinction matters.

[7]Kagan (1989) denies that a convincing moral theory must acknowledge agent-relative reasons. But this is a minority position.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

What exactly makes a reason agent-relative? On a common understanding, a reason is agent-relative if its full specification includes an ineliminable pronominal back-reference to the agent for whom it is a reason, and it is agent-neutral otherwise (Nagel 1986; McNaughton and Rawling 1998; Darwall 2006). It is common to distinguish between three kinds of agent-relative reasons, namely project-dependent reasons, relationship-dependent reasons, and restrictions (Nagel 1986).[8] Project-dependent reasons are reasons that stem from the personal projects of agents: Arnold has a special reason to pursue *his* project. Relationship-dependent reasons stem from the personal relationships of agents: Bernard has a special reason to benefit *his* loved one or express *his* commitment to a relationship. And restrictions are reasons not to perform acts of certain types: Christine has a special reason not to lie *herself*, rather than prevent others from lying. There might be other kinds of agent-relative reasons,[9] but these three kinds are widely recognized as paradigm cases, and I will focus on them in what follows.

There is some disagreement in the literature on how to interpret the reference to the agent. Some authors understand it as a motivational feature (Pettit 1987, 1988; Double 1999). On this interpretation, the fact that *my* daughter is drowning motivates me to jump into the water to save her. This interpretation is incomplete at best: the reference to the agent certainly plays a normative role and (partly) explains why an agent ought to act in a certain way. Most people think that a father who saves a stranger from drowning rather than his daughter does not merely exhibit a motivational shortcoming but fails to act on a normative reason to save his daughter and acts wrongly. Hence, the reference to the agent seems to be relevant in a normative sense and, again, this can mean several things. It can mean that the reference to the agent is normatively relevant because it generates a reason, and that the fact that *my* daughter is drowning gives me an agent-relative reason to save her. A reason is therefore agent-relative if it is a reason only for the agent but not for everybody (Raz 1986; Scheffler 2004; Jeske 2008; Keller 2013). But it can also mean that the reference to the agent affects the normative force of an already existing reason (van Willigenburg 2005; Wallace 2010): the fact that a girl is drowning gives me a reason to save her, but the fact that this girl is *my* daughter intensifies the force of that reason for me. On the first view, a reason is agent-relative if it exists only for the agent who has it: *I* have an agent-relative reason to save *my* daughter that Frank, who is not her father, does not have (even though Frank has an agent-neutral reason to save her). On the second interpretation, a reason is agent-relative if it has different authority over different agents: Frank and I both have the same reason to save my daughter (because we both have reason to save a drowning child), but my reason to save that particular child is stronger than Frank's.[10]

Both interpretations of the agent-relative/agent-neutral distinction pose a similar problem for AVR. According to AVR, practical reasons are reasons for appropriate value responses. Proponents of AVR should therefore accept something like the following account of agent-relative reasons:

> *Agent-Relative Reasons (AVR)*: A reason is agent-relative if its constituting an appropriate response to a certain agent-neutral value *V* depends on the identity of the responding agent, and it is agent-neutral otherwise.[11]

So, for example, I have an agent-relative reason to risk my life in order to save my daughter from drowning or to throw a life preserver to my daughter rather than to two children of a stranger

[8]Nagel calls project-dependent reasons "reasons of autonomy" and relationship-dependent reasons "reasons of obligation."

[9]Dancy (1993) identifies seven kinds of agent-relative reasons.

[10]Depending on one's metaphysical views about reasons, these two interpretations need not come apart. Some authors think that there is such a thing as a complete reason which consists of all the normatively relevant features of a situation (Raz 2000; Heuer 2006). If the features that determine the strength of a reason are among this set of relevant considerations, then the strength of the same reason cannot differ in two situations because two reasons that differ in strength are two different (complete) reasons.

[11]Thanks to Tristram McPherson for suggesting this formulation to me.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

because these responses to the value of my daughter are appropriate for me. But they might not be appropriate responses for strangers and, in this sense, I have agent-relative reason to do these things that other people do not have.

Note that the notion of appropriateness is not an overall notion here. The issue is not whether the appropriateness of an overall response to a situation depends on the identity of the agent, but whether the appropriateness of a response to a certain value does. This is important for the account to yield the right results. For example, it leaves open the possibility that an action is supported both by agent-relative and agent-neutral reasons.[12] This can happen when an agent can respond to more than one value at once by performing some action. Suppose, for example, that my only options right now are to press a button or refrain from doing so.[13] If I press the button, millions of strangers will be saved from disaster. This is intuitively an agent-neutral reason to press the button. Suppose furthermore that I enjoy pressing the button, which is intuitively an agent-relative reason for me to press the button. In such a case, the appropriateness of me pressing the button does not depend on the latter reason since the former is so strong; and it is certainly appropriate for everybody to press the button. Hence, the appropriate overall response to the situation (namely pressing the button) does not differ from agent to agent and therefore does not depend on my identity in particular. Nevertheless, I do have an agent-relative reason to press the button since pressing the button is also an appropriate response to the value of the pleasure that I get from doing so. Hence, in this case, one and the same overall response is supported by an agent-relative as well as an agent-neutral reason. And, of course, we can also imagine cases in which an agent has agent-neutral reason to do one thing, but agent-relative reason to do another—it all depends on the agent-neutral values that the agent can respond to in a given situation. And this is how it should be: every plausible account of agent-relative and agent-neutral reasons should allow for the possibility of such reasons to conflict.

Understanding agent-relative reasons in terms of appropriate responses to specific agent-neutral values rather than in terms of the appropriateness of overall responses is also important because otherwise the following problem could arise. As an anonymous referee pointed out to me, it is only possible to compare the appropriateness of different agents performing the same action when we think of the action as a type, and once we describe the act of saving one's daughter as an instance of the act type "daughter saving" (which by definition one only performs by rescuing his or her own daughter), then the account might let my reason to save my daughter appear agent-neutral. After all, it is appropriate for everyone to save their own daughter, and the appropriateness of my response therefore does not seem to depend on my specific identity. But once we understand agent-relative reasons in terms of appropriate responses to specific agent-neutral values and not in terms of overall responses, this problem does not arise. The question is what response to the value of *this specific child* is appropriate for me. And as a response to this specific value, it is appropriate for me to do different things than is appropriate for strangers to do.

We can now see what problem agent-relative reasons pose for a value-based theory of reasons. If all value is agent-neutral value, then it seems natural to suppose that the appropriateness of a value response is determined by V, not by the responding agent, and the appropriateness of a value response should not vary across agents. Hence, if the drowning of my daughter is agent-neutrally bad, then everybody should have the same reason to respond to this (dis-)value. But agent-relative reasons suggest otherwise. The question that proponents of a value-based theory of practical reasons must answer is how it can be appropriate for different agents to respond differently to the same agent-neutral value.

In what follows, I will set out to answer this question and offer a novel account of accommodating agent-relative reasons within an AVR approach that understands value as agent-neutral. This account will cover all kinds of agent-relative reasons and show that they have a unified

[12]I thank an anonymous referee for pressing me on this point.

[13]I thank Tristram McPherson for this example.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

rationale, which I take to be an advantage of the account.[14] More specifically, I will explain agent-relative reasons in terms of higher-order value responses. Let me explain.

4. Agent-relative reasons as second-order value responses

There are two broad strategies to accommodate agent-relative reasons within AVR. First, one can start with some agent-neutrally valuable entity and then go on to explain why it is appropriate for different agents to respond differently to this entity. Proponents of this strategy therefore start with the plausible claim that the survival of my daughter is good and then go on to try to explain why different responses to my daughter's survival are appropriate for you and for me respectively. The underlying assumption is that when my daughter is drowning and I jump into the water, while you throw a life preserver into the water and call for help, we both respond to the same valuable entity (the value of my daughter, or perhaps the value of the state of affairs of my daughter surviving), but different responses are appropriate for us. Many authors endorse this strategy, and while I do not have the space here to discuss all variations of this strategy, it seems safe to say that they all face considerable objections.[15] Hence, I will not pursue this first strategy here. Rather, I will pursue a second strategy to accommodate agent-relative reasons within AVR and specify which valuable entities agents respond to when they act on their agent-relative reasons. My claim will be that when you and I both try to save my daughter, we both respond to agent-neutrally valuable entities, but to different ones. And since we respond to different entities, we have different reasons. To demonstrate this, I start with project-dependent reasons and discuss them in some detail. Subsequently, I show that the results of this discussion generalize to relationship-dependent reasons and restrictions.

4.a Project-dependent reasons

Which entity does an agent respond to when he acts on a project-dependent reason? An obvious answer seems to be that he responds to the value of the end the project is based on. For example, if my project is to write a book on agent-relative reasons, and if this end is genuinely valuable, then I respond to the value of this end by spending a weekend outlining a chapter of the book. And if your project is to run a marathon, and if that end is genuinely valuable, then you respond to the value of that end by spending the weekend training for stamina. Call this the *straightforward account* of project-dependent reasons.

This straightforward account has several advantages. It explains the intuition that agents have no reason to pursue immoral projects. For instance, according to the straightforward account, a member of the Ku Klux Klan has no reason to burn churches because his project is not based on a valuable end. Furthermore, the account is in line with (at least some aspects of) the phenomenology of pursuing projects. We often take ourselves to respond to the value of our project when we pursue it. If you are fed up with training for the marathon, you might gain new motivation by focusing on what an achievement it will be to cross the finish line.

Nevertheless, I think that the straightforward account falls short. First, it does not really account for agent-relative reasons. If completing a marathon and writing a book on agent-relative reasons are both valuable ends, then I seem to have just as much reason to train for stamina as I have reason

[14]Not all authors might agree: some authors accept project- and relationship-dependent reasons but reject restrictions (Scheffler 1982; Jeske 2001; McNaughton and Rawling 2013). However, since one important motivation to accommodate agent-relative reasons within a moral theory is that they are intuitively compelling, and since restrictions are intuitively compelling, it seems desirable to develop an account that includes restrictions.

[15]It underlies the many attempts to accommodate agent-relative reasons within a consequentialist framework, such as rule-consequentialism, indirect consequentialism, motive consequentialism, etc. It is also the background of Barry Maguire's (2017) recent attempt to reconcile agent-relative reasons with an agent-neutral value-first view. And it can be found in nonconsequentialist attempts to explain partiality within a value-based view, such as Oddie (2005).

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

to work on a chapter for the book. According to the straightforward account, what is relevant is whether the end is valuable, not whether it is my project. Thus, the question remains why the one valuable end gives me agent-relative reasons for action, whereas the other does not. Furthermore, why should I have reason to pursue my project *myself*, rather than to help others to achieve the valuable goal? Suppose that someone else can write a much better book on agent-relative reasons than me. Why should it be an appropriate response to the value of this end to try to write the book myself? It would seem more appropriate to help the other person write the better book. Second, phenomenological considerations not only support the straightforward account but also speak against it. We often regard the fact that some project is ours as equally reason-providing as the fact that the project aims at a valuable end (Betzler 2004). If I am frustrated with the slow progress of my book and cease to consider finishing it a valuable achievement, I might gain new motivation by thinking about the time and energy that I have already invested in the project. Third, the straightforward account has problems accounting for projects that are based on unrealizable ends. Suppose that my project is to help keep global warming below 2°C. At some point, I might realize that this goal is unachievable. Yet, it does not seem unreasonable for me to carry on with my project and educate people about the dangers of global warming, write petitions to politicians, and so on. This is difficult to explain if project-dependent reasons were based solely on the value of a project's end, given that this value will never be realized.

A better account of project-dependent reasons starts from a different understanding of the value of a project. When we talk about the value of a project, we can either refer to the valuable end the project is based on or we can refer to the value of pursuing the project. It is this latter kind of value that agents respond to when they act on their project-dependent reasons.

Consider what it means to respond appropriately to some value *V*. One possible way of responding appropriately to *V* is to maximize *V*, but there are also other candidates—take, for example, Joseph Raz's (2001) distinction between recognizing value, respecting value, and engaging with value. Recognizing value involves "regarding objects in ways consistent with their value" (161). Respecting value involves the stance "not to destroy, and furthermore, to preserve what is of value" (162). Engaging with value involves making the value a part of one's life, making oneself emotionally vulnerable to the value, etc. According to Raz, this latter kind of value response is the most appropriate one, since value is ultimately realized when it is engaged with (163). It is also the most demanding one, since it requires agents to make use of their limited resources (time, money, energy, etc.).

Pursuing a project is a way of engaging with value. Once I adopt the project of writing a book on agent-relative reasons, I engage with the value of that intellectual achievement. I make this value part of my life, and I make myself emotionally vulnerable to it: I am prone to feel frustrated if my work is not going well, to feel proud if I have managed to write a chapter that satisfies my standards, and so on. When an agent pursues a project, she engages with value—namely with the value of the project's end—and she therefore responds to value in the most appropriate form.

As a way of engaging with value, pursuing a project is intrinsically valuable. This intrinsic value can be brought out in two different ways. First, as Thomas Hurka (2001) has convincingly argued, appropriate responses to value—whatever they are—are intrinsically good. Compare two worlds, W1 and W2. W1 contains both suffering and no compassion, whereas W2 contains the same amount of suffering as well as compassion. Intuitively, W2 seems to be the better world due to the compassion. Compassion is an appropriate response to the (dis-)value of suffering, and since the mere occurrence of compassion in W2 makes it better independent of any further consequences (the amount of suffering is the same), this suggests that it is intrinsically good to respond appropriately to value. Of course, compassion is an attitude, not an action, but it seems unproblematic to extend this claim to actions: acting in a way that is an appropriate response to value is intrinsically good, independently of any other consequences that the action might bring about. Second, by adopting a project, an agent structures (parts of) her life and brings (at least some of) her desires, intentions, emotions, etc. into a coherent whole. She therefore unifies quite disparate

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

elements. And according to an attractive metaphysical view of value, value supervenes on the unification of disparate elements into complex wholes (Nozick 1981; Oddie 2005; Kelly 2014). It is better to live a structured life than to live an unstructured life, and since pursuing a project consists in structuring (at least parts of) one's life, this gives an independent explanation of why it is valuable to pursue projects. Taken together, these two considerations suggest that pursuing projects is intrinsically valuable, and that there is value in projects over and above the value of their ends.

If pursuing a project is intrinsically valuable, then there is also reason to respond appropriately to this value. Such appropriate responses include attitudes as well as actions. An agent responds appropriately to the value of pursuing a project by being proud of the project or by hoping that she will carry it out successfully, and she also responds appropriately to the value of pursuing a project by performing actions that help facilitate the project's success. In other words, she responds appropriately to the value of pursuing a project by acting on her project-dependent reasons. Since pursuing a project is a way of engaging with value and is therefore itself an appropriate value response—namely a response to a valuable end—acting on a project-dependent reason is an appropriate response to an appropriate value response. In other words, acting on a project-dependent reason is a second-order value response.

Understanding agent-relative reasons in terms of second-order value responses accommodates them within an AVR framework. Agent-relative reasons do not arise first and foremost because it is appropriate for different agents to respond differently to the same agent-neutrally valuable entities; they arise because it can be appropriate for different agents to respond to different agent-neutrally valuable entities in a situation. You and I have different project-dependent reasons because we have different first-order value responses and we have reason to respond to our different first-order value responses.[16]

Some readers might find this account implausible. I have claimed that practical reasons are reasons to respond appropriately to agent-*neutral* value, but one might think that the value of a first-order value response is agent-*relative* value since it is *the agent's* first-order value response. However, the mere fact that it is the agent's response does not make the value agent-relative. The value is agent-*relational*: its existence depends on valuing agents.[17] But that does not mean that it is also agent-*relative* in the sense that its good-making features include reference to the agent. My appropriate value response exists because of me, but the fact that my appropriate value response is good does not depend on me. The value of an appropriate value response is agent-relational, yet agent-neutral, value.

Other readers might object that, if pursuing projects adds value to a person's life, then this means that it is good *for* a person to pursue projects, and that the "good for"-relation expresses agent-relative value. However, independently of the question whether the "good for"-relation does in fact express agent-relative value,[18] the fact that pursuing projects adds value to my life does not entail that the structuring of my life is good for me. After all, I can adopt extremely self-denying projects. In such a case, my life is structured, and this adds value to my life, but it is not necessarily good for me. Hence, even if the "good for"-relation expresses agent-relative value, it does not follow

[16]An editor of this journal has pressed me to explain why invoking the higher-order level of value response is necessary to explain agent-relative reasons. Doesn't Raz already solve the puzzle insofar as the appropriateness of the engagement response already accounts for agent-relative reasons, so that the fact that doing *A* presents a way for *S* to engage with an agent-neutral value *V* is an agent-relative reason for *S* to do *A*? While my account is close to this idea, I think that it is not enough to rely merely on the engagement response to explain agent-relative reasons. Most importantly, relying merely on the engagement response does not explain deontological restrictions because respecting a restriction is not a way of engaging with value. By contrast, my account gives a unified explanation of project-dependent reasons, relationship-dependent reasons, and restrictions. I will return to this below.

[17]I take the term "agent-relational" from Schroeder (2007b). However, Schroeder talks about agent-relational reasons, not agent-relational value.

[18]Much depends on how to understand the "good for"-relation. This is a contested matter, and I cannot discuss it here. For different accounts, see Hurka (1987), Regan (2004), Rosati (2008), and Rønnow-Rasmussen (2011).

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

that the value of pursuing personal projects is agent-relative because pursuing projects is not necessarily good for the agent.

Still, some readers might insist that even if pursuing projects is not necessarily good for an agent, the mere fact that it adds value to a life makes the value agent-relative. But this does not follow either. Suppose that the Great Consequentialist Project has succeeded, and we have made the world as good as possible. This seems to add value to everybody's life even though it is not necessarily good for everybody (it allows for trade-offs between people): *ceteris paribus*, living in the best possible world is better than living in a world that is not quite as good. But this does not show that the success of the Great Consequentialist Project has agent-relative value. On the contrary, it is a paradigmatic case of agent-neutral value.

I conclude that pursuing personal projects is agent-neutrally good. For its good-making features do not depend on the identity of the agent, and its value is neither agent-relative in the sense that it is necessarily good for the agent to pursue projects, nor is its value agent-relative in the sense that pursuing projects adds value to a life. Since the value of pursuing a project generates agent-relative project-dependent reasons, and since the value of a project is agent-neutral value, it follows that agent-neutral value generates agent-relative reasons.

Two open questions remain at this point. First, I have criticized the straightforward account of project-dependent reasons because it implies that I have just as much reason to train for stamina as I have reason to write a chapter for my book. What remains to be shown is that understanding agent-relative reasons in terms of second-order value responses fares better in this regard. Second, a decisive feature of agent-relative reasons is that they are reasons to take one's own projects, relationships, or integrity as more important than the projects, relationships, or integrity of others, even though the projects, relationships, and integrity of others are equally agent-neutrally valuable.[19] What remains to be shown is that understanding agent-relative reasons in terms of second-order value responses explains this feature. I will discuss these points in turn.

Regarding the first point, the important difference between the straightforward account and the view proposed here is how the respective accounts specify what values agents respond to when they act on their project-dependent reasons. On the straightforward account, I have just as much reason to work on a chapter of my book as I have reason to train for stamina because in each case, I respond directly to an agent-neutral base-level value. On the view proposed here, this is not the case. When I work on the chapter for my book, I do not simply respond to the base-level value of that intellectual achievement, but to my own response to the base-level value. This explains why other agents who do not have the same first-order value response to the base-level value do not have the same reason as me: they cannot respond to the value of that specific first-order value response.

Of course, I *could* adopt the project of completing a marathon, and one could argue that the value of the athletic achievement of running a marathon gives every agent a reason to adopt a corresponding project, and that therefore every agent has the same reason to train for stamina after all. However, it is important to remember that pursuing projects is a way of engaging with value. Since engaging with value makes significant demands on agents and their resources (time, money, energy, etc.), nobody can engage with any possible value. Agents must choose with which values they engage. They have leeway in choosing their projects and they have reason to choose any project that is based on a valuable end. Hence, if both running a marathon and writing a book on agent-relative reasons are valuable ends, I have *ceteris paribus* just as much reason to choose the one project as I have reason to choose the other. However, once I have chosen a value with which to engage and have adopted a corresponding project, I need to invest my time, energy, and other resources in a way that makes it impossible to adopt other projects in the same manner.

This brings me to the second point: the special normative force of agent-relative reasons. Whenever an agent acts on her agent-relative reasons, it is important that the agent *herself* performs

[19]I thank an anonymous referee for helping me spell out this problem better than in an earlier version of the paper.

some action or refrains from doing it, rather than helping others perform actions of that very type.[20] For example, I have an agent-relative reason to facilitate the success of my project rather than helping some other person pursue her project. I also have an agent-relative reason to take care of my child even if, by neglecting my own child, I could bring another parent to take care of his or her child who would otherwise neglect the child. And I have a reason not to tell a lie even if my lie could prevent someone else from lying. This special significance of agent-relative reasons is the litmus test for deciding whether a reason is agent-relative or not (Darwall 2006) and any account of agent-relative reasons must be able to make sense of it.

How can the proposed view account for this special significance of agent-relative reasons? The first thing to remember is that an agent who acts on her project-dependent reasons responds not merely to some value, but more specifically to the value of her own appropriate value response. This value only exists because she responds to value in the way that she does, and it includes structuring her life accordingly. She can only structure her life if she does *not* take herself to have just as much reason to respond to her own project as she has to respond to the project of others. If an agent does not take herself to have special reason to respond to her own project, then she does not structure her life around the project's end. After all, she always takes it as a salient possibility to do something else other than bringing her own actions, emotions, etc. into a coherent whole. And this means that the agent could not realize the value of her own value response if she does not take herself to have special reason to respond to her own value response rather than the value responses of others.

Of course, an agent also has reason to respond to the agent-neutral value of someone else's value response. But that other person's value response also exists on the condition that the other person takes herself to have special reason to pursue her own project rather than helping others pursue their projects. In other words, my value response only exists because I take myself to have special reason to pursue my own project; the other person's value response exists because she takes herself to have special reason to pursue her project. Thus, even if I respond to another person's value response, the value to which I respond depends on the general principle that agents have special reason to take their own projects as more important than the projects of others. This means that by responding to the other person's project, I affirm that appropriate value responses are proper objects of further value responses. And since such appropriate value responses are only possible if agents take themselves to have special reason to respond to their own value responses rather than the value responses of others, I indirectly affirm that I have special reason to respond to my own project even if I respond to someone else's value response.

We can couch the underlying principle in terms of appropriateness. An appropriate response to a value is a pro-response, and that means that an appropriate response to a value should not make the very value it responds to impossible. But this would be the case if it were just as appropriate for an agent to respond to his own value response as it were for her to respond to the value responses of others. To say that an agent has as much reason to help others act on their project-dependent reasons as she has to act on her own means that it would be just as appropriate to do the one thing as the other. Since helping others to act on their project-dependent reasons implies not being able to perform acts that are part of the interconnected actions that constitute the structuring of her own life, this means that it would be just as appropriate to respond to one's own value response as it would be to act in a way that makes one's own value response impossible. Of course, the value that is potentially made impossible and the value that is responded to are not the same values: the former is my value response and the latter is the value response of someone else. However, for the other person it would also be just as appropriate to forsake her own value response for the sake of the value

[20]Many authors discuss this special normative force of agent-relative reasons in terms of nonmaximization: I have agent-relative reason to φ even if I could maximize or minimize actions of the very same type by not φ-ing. However, the special normative force of agent-relative reasons already occurs when maximization is not at issue. I thank an anonymous referee for helping me to clarify this.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

responses of others, and she could therefore also not structure her life around a value. In sum, if it were just as appropriate to respond to my own value response as it would be to respond to the value responses of others, then this would mean that an appropriate value response could threaten the very same value that it is a response to. It is therefore more appropriate to focus on one's own value response than it is to respond to the value responses of others.

The special normative significance of project-dependent reasons can also be explained another way. When an agent structures her life around a value, then her actions at t gain much of their meaning from her actions at t_1, t_2, etc. Her reasons to perform various acts are thus intertwined, and every instance in which she acts on a project-dependent reason has significance beyond itself. The normative force of her project-dependent reasons therefore also extends beyond the individual situation in which a specific reason arises. By contrast, the normative force of her reason to help someone else pursue his project does not extend beyond itself. Hence, the agent's projects have a normative significance for her that the projects of others do not.

To conclude, project-dependent reasons can be accommodated within an agent-neutral AVR framework by understanding them in terms of second-order value responses. What remains to be shown is that the account generalizes to other kinds of agent-relative reasons as well.

4.b Relationship-dependent reasons

It seems unproblematic to extend the account to relationship-dependent reasons. Personal relationships are like projects in many ways. They too are ways of engaging with value, namely the value of persons (Velleman 1999; Rosati 2008; Keller 2013). Furthermore, they structure the lives of their participants in a similar way as do projects. In fact, relationships constitute complex unity to an even greater degree. Not only do they structure the plans, actions, and desires of an individual agent and bring them into a coherent whole, they also bring plans, actions, and desires of two individual agents into one coherent whole. Relationships are therefore intrinsically good in the same way projects are, and this means that they are also objects of appropriate value responses. And one way to respond appropriately to the value of one's relationship is to act on one's relationship-dependent reasons and spend time with one's friend, or care for one's child. Thus, the explanation of project-dependent reasons applies to relationship-dependent reasons as well, and so does the explanation of their special normativity.

Understanding relationship-dependent reasons in terms of second-order value responses has an important further advantage. It helps resolve a prominent debate in the ethics of personal relationships. Some authors claim that relationship-dependent reasons are grounded in the value of a relationship (Jeske 2008; Tan 2010; Seglow 2013). Others argue that such reasons are grounded in the value of the person that we have the relationship with (Keller 2013). On the view proposed here, there is truth in both views. When we act on a relationship-dependent reason, we respond appropriately to the value of a relationship and, in this sense, the reason stems from the value of the relationship. However, the relationship is valuable (and thus reason-giving) because it is an appropriate response to the value of the other person and, in this sense, the ultimate ground of the relationship-dependent reason is the value of the other person. Thus, when we act on a relationship-dependent reason, we respond directly to the value of a relationship but we also respond indirectly to the value of the other person.

4.c Restrictions

Deontological restrictions do not obviously fit the model of project- and relationship-dependent reasons. Neither is respecting a restriction a way of engaging with value, nor do restrictions structure the lives of the agents who respect them. Nevertheless, restrictions can also be interpreted in terms of second-order value responses. To show this, it is again important to ask which value an agent responds to when she respects a restriction.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

One possible answer is that restrictions are appropriate responses to *dis*values, and that by refraining from lying or killing, an agent responds appropriately to the disvalue of lying or killing. But this answer does not explain the agent-relativity of restrictions: refraining from telling a lie and preventing someone else from telling a lie would be responses to the very same value, and I would have just as much reason to do the one as the other even if the latter includes telling a lie myself. This first way of explaining restrictions lets them even appear paradoxical (Scheffler 1982; McNaughton and Rawling 2013). If restrictions are grounded in the disvalue of certain act types, then it seems paradoxical to hold that an agent should respect a restriction even if she could minimize the occurrence of such act types by violating a restriction herself. Rather, she should lie when she can thereby prevent even more agents from lying. But this is not how we conceive of restrictions.

Another possibility is that restrictions are appropriate responses to the value of persons, and that refraining from torturing a person or from lying to her are appropriate responses to her value. However, this view leaves restrictions somewhat unmotivated. Why do I respond appropriately to the value of a person by refraining from torturing her rather than by preventing her from being tortured? The value of a person surely generates reasons to help her avoid mistreatment, and thus both respecting restrictions as well as violating them in order to prevent violations by others could be appropriate responses to the value of persons.[21]

A third and better approach understands restrictions in terms of second-order value responses. Agents do not respond appropriately to the value of persons merely by respecting restrictions individually. Rather they do so by establishing social relations that include deontological rules. When agents collectively create social relations with deontological rules, they thereby collectively respond appropriately to their respective value as free and equals. Deontological restrictions protect the agency of persons, and this enables them to form and pursue life plans. In this sense, restrictions are appropriate responses to the value of persons as free (or autonomous) beings. And since deontological rules are universal rules in the sense that they equally protect all members of the moral community, agents also express their status as equals by creating social relations with such rules.[22] Thus, social relations that include deontological rules are similar to projects and relationships in that they are also an appropriate response to value. And they share another similarity. Respecting restrictions might not count as engaging with value, and restrictions might not structure the lives of individual agents, but they structure the actions and intentions of different agents insofar as what is permissible for me to do influences what is permissible for you to do and vice versa. Therefore, social relations that include restrictions exemplify at least some degree of unified complexity.

Since social relations with deontological rules are an appropriate response to value, they are intrinsically good and therefore objects of appropriate value responses. And agents respond appropriately to the value of such social relations by respecting the deontological rules that govern them. This explains why an agent does *not* have as much reason to respect a restriction herself as she has reason to prevent others from violating a restriction, and it explains why deontological restrictions are not paradoxical.[23] Restrictions only appear paradoxical if one understands them as responses to the disvalue of certain act types. Understanding restrictions in terms of second-order value responses rejects this assumption. One might argue that if restrictions are appropriate responses to the value of a certain kind of social relations, and if violating restrictions is intrinsically bad in virtue of being an inappropriate response to this value, then the appropriate response to the value of such social relations is to maximize compliance with the deontological rules that govern

[21]See Lippert-Rasmussen (2009) who argues that the moral status of a person can be described as "inviolability," but also as "unignorability." In the latter case, agents must not ignore potential victims of the violation of restrictions, which includes those persons who the agent could prevent from becoming victims by violating a restriction herself.

[22]By claiming that the value of a person is the base-level value to which agents respond when they respect a deontological restriction, my account is similar to Frances Kamm's (1992, 2007) or Thomas Nagel's (1995) accounts of deontology. I thank an anonymous referee for pointing out the similarities between my account and Kamm's.

[23]Other answers to the paradox of deontology include Mack (2005), Heuer (2011), and Chappell (2014).

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

them. But at this point, a similar response to the one available in the case of projects and relationships is available here. It is only possible to realize the value of social relations that include deontological rules if agents are *not* taken to have as much reason to respect restrictions as they have preventing others from violating them. If agents had as much reason to prevent someone from lying as they have reason not to lie themselves (and if they would therefore be morally permitted to lie in order to prevent someone else from lying), it could hardly be said that their social relations are governed by deontological rules. Once we accept that restrictions are responses to the value of certain kinds of social relations rather than responses to the disvalue of certain act types, we can account for their special agent-relative normative force.

Another advantage of the view developed here is that it allows for the possibility of a threshold deontology.[24] Any plausible deontology must permit or even require agents to violate a restriction if this is the only way to prevent catastrophic disasters. Understanding restrictions in terms of higher-order value responses accommodates this. Agents are not permitted to tell one lie in order to prevent two lies from being told, but they are permitted—or even required—to tell a lie if that is the only way to prevent the whole socially embodied set of norms that constitutes the common appropriate response from collapsing. I cannot spell out this possible threshold deontology in detail here, but it is worth noting that the view proposed allows for such a threshold deontology, which I take to be an advantage of the view.

5. Conclusion

An agent-neutral value-first framework can accommodate agent-relative reasons. The important insight is that appropriate value responses can occur on different levels, and that appropriate responses to agent-neutral values are themselves of agent-neutral value. Despite important differences between project-related reasons, relationship-dependent reasons, and restrictions, there is a unifying framework which explains them. They all can be described in terms of second-order value responses—that is, in terms of appropriate responses to appropriate value responses.

In closing, I would like to mention some advantages of the proposed view. First, it meets criteria that a convincing account of agent-relativity should meet: it covers different kinds of agent-relative reasons and it explains why agents ought to act on their own agent-relative reasons rather than helping others acting on theirs (or maximizing general compliance with such reasons). Second, the view nicely accounts for the two-sided phenomenology of acting on agent-relative reasons. When pursuing a project, we might take ourselves to respond directly to the valuable goal on which the project is based, or we might take ourselves to respond to the fact that we are pursuing this specific project. Similarly, I might help my friend out of concern for her or out of concern for our friendship. And I might respect a restriction out of respect for the other person or out of concern for the moral rules that ground the restriction. All these possibilities seem adequate and the view developed here explains why. The actual object of response is the agent's first-order value response and it is adequate to focus on the actual object of one's response. But since the base-level value grounds the first-order value response, it can also be adequate to focus on the ground of the actual object of one's response—in a sense, that value shimmers through.

Another advantage of this view is that it does not understand agent-relative reasons and agent-neutral reasons as incommensurable. Since they are both ultimately grounded in agent-neutral value, it should (at least in principle) be possible to weigh them against each other and determine what an agent ought to do, all things considered, when they conflict. This is an advantage because if agent-relative reasons and agent-neutral reasons were incommensurable, then every case in which an agent-neutral reason conflicts with an agent-relative reason would constitute a dilemma, and that is hardly convincing. I cannot discuss here exactly how this weighing works; it depends on how

[24]I thank an anonymous referee for pointing this out to me.

https://doi.org/10.1017/can.2020.3 Published online by Cambridge University Press

values can be compared in general and how AVR accounts for the strength of practical reasons. The important point is that we should not assume that agent-relative reasons and agent-neutral reasons are incommensurable merely because the former involve an essential reference to the agent while the latter do not.

Of course, explaining agent-relative reasons in terms of second-order value responses is not the only way to account for such reasons. What I hope to have shown, however, is that agent-relative reasons pose no problem for a value-based theory of practical reasons that understands all value as agent-neutral. Even if all value is agent-neutral, agent-relative reasons emerge naturally from what it means to respond appropriately to value.

Acknowledgments. For helpful comments, I thank audiences in Zurich, Bern, Munich, Southampton, and Cologne, especially Monika Betzler, Christoph Halbig, Tristram McPherson, and Thomas Hurka, as well as Andreas Blaser, an editor of this journal, and three very helpful referees.

Funding Statement. 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).

Jörg Löschke is currently an SNSF-funded professor at the University of Zurich in the department of philosophy. His research focus is on value theory and normative ethics, especially the distinction between consequentialist and nonconsequentialist moral theories and the ethics of personal relationships.

References

Betzler, Monika. 2004. "Personal Projects as Practical Reasons." Unpublished manuscript.
Bykvist, Krister. 2009. "No Good Fit: Why the Fitting Attitude Analysis of Value Fails." *Mind* 118: 1–30.
Chappell, Timothy. 2014. *Knowing What to Do. Imagination, Virtue, and Platonism in Ethics*. Oxford: Oxford University Press.
Crisp, Roger. 2005. "Values, Reasons and the Structure of Justification: How to Avoid Passing the Buck." *Analysis* 65: 80–85.
Dancy, Jonathan. 1993. *Moral Reasons*. Oxford: Blackwell.
Dancy, Jonathan. 2000. "Should We Pass the Buck?" *Royal Institute of Philosophy Supplement* 47: 159–73.
Danielsson, Sven, and Jonas Olson. 2007. "Brentano and the Buck-Passers." *Mind* 116: 511–22.
Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
Double, Richard. 1999. "Morality, Impartiality, and What We Can Ask of Persons." *American Philosophical Quarterly* 36: 149–58
Finlay, Steven. 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
Gregory, Alex. 2014. "A Very Good Reason to Reject the Buck-Passing Account." *Australasian Journal of Philosophy* 92: 287–303.
Gregory, Alex. 2016. "Reasons as Good Bases." *Philosophical Studies* 173: 2291–310.
Hammerton, Matthew. 2016. "Patient-Relativity in Morality." *Ethics* 127: 6–26.
Heuer, Ulrike. 2006. "Explaining Reasons: Where Does the Buck Stop?" *Journal of Ethics & Social Philosophy* 1: 1–25.
Heuer, Ulrike. 2011. "The Paradox of Deontology, Revisited." *Oxford Studies in Normative Ethics* 1: 236–67.
Hurka, Thomas. 1987. "'Good' and 'Good for.'" *Mind* 96: 71–73.
Hurka, Thomas. 2001. *Virtue, Vice, and Value*. Oxford: Oxford University Press.
Hurka, Thomas. 2003. "Moore in the Middle." *Ethics* 113: 599–628.
Hurka, Thomas. 2014. *British Ethical Theorists from Sidgwick to Ewing*. Oxford: Oxford University Press.
Jeske, Diane. 2001. "Friendship and Reasons of Intimacy." *Philosophy and Phenomenological Research* 63: 329–46.
Jeske, Diane. 2008. *Rationality and Moral Theory: How Intimacy Generates Reasons*. New York: Routledge.
Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Clarendon Press.
Kamm, Frances. 1992. Non-Consequentialism, the Person as an End-in-itself, and the Significance of Status. *Philosophy and Public Affairs* 21: 354–89.
Kamm, Frances. 2007. *Intricate Ethics*. Oxford: Oxford University Press.
Kauppinen, Antti. 2014. "Fittingness and Idealization." *Ethics* 124: 572–88.
Keller, Simon. 2013. *Partiality*. Princeton, NJ: Princeton University Press.
Kelly, Chris. 2014. "Value Monism, Richness, and Environmental Ethics." *Les ateliers de l'éthique/The Ethics Forum* 9: 110–29.
Lang, Gerald. 2008. "The Right Kind of Solution to the Wrong Kind of Reason Problem." *Utilitas* 20: 472–89.
Lippert-Rasmussen, Kasper. 2009. "Kamm on Inviolability and Agent-Relative Restrictions." *Res Publica* 15: 165–78.
Louise, Jennie. 2004. "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518–36.
Mack, Eric. 2005. "Prerogatives, Restrictions, and Rights." *Social Philosophy and Policy* 22: 357–93.
Maguire, Barry. 2016. "The Value-Based Theory of Reasons." *Ergo* 3: 232–62.

Maguire, Barry. 2017. "Love in the Time of Consequentialism." *Noûs* 51: 686–712.
McNaughton, David, and Piers Rawling. 1992. "Honoring and Promoting Values." *Ethics* 102: 835–43.
McNaughton, David, and Piers Rawling. 1998. "On Defending Deontology." *Ratio* 11: 37–54.
McNaughton, David, and Piers Rawling. 2013. "Contours of the Practical Landscape." In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, edited by David Bakhurst, Brad Hooker, and Olivia Little, 240–64. Oxford: Oxford University Press.
McHugh, Conor, and Jonathan Way. 2016. "Fittingness First." *Ethics* 126: 575–606.
Nagel, Thomas. 1986. *The View from Nowhere*. New York: Oxford University Press.
Nagel, Thomas. 1995. "Personal Rights and Public Space." *Philosophy and Public Affairs* 24: 83–107.
Nair, Shyam. 2014. "A Fault Line in Ethical Theory: Consequentialism, Deontic Constraints, and the Prisoner's Dilemma." *Philosophical Perspectives* 28: 173–200.
Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
Nye, Howard, David Plunkett, and John Ku. 2015. "Non-Consequentialism Demystified." *Philosophers' Imprint* 15: 1–28.
Oddie, Graham. 2005. *Value, Reality, and Desire*. Oxford: Oxford University Press.
Orsi, Francesco. 2013. "What's Wrong with Moorean Buck-Passing?" *Philosophical Studies* 164: 727–46.
Pettit, Philipp. 1987. "Universalizability without Utilitarianism." *Mind* 96: 74–82.
Pettit, Philipp. 1988. "The Paradox of Loyalty." *American Philosophical Quarterly* 25: 163–71.
Portmore, Douglas. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. "The Strike of the Demon: On Fitting Pro-Attitudes and Values." *Ethics* 114: 391–423.
Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2006. "Buck-Passing and the Right Kind of Reasons." *Philosophical Quarterly* 56: 114–20.
Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Clarendon Press.
Raz, Joseph. 2000. "The Truth in Particularism." In *Moral Particularism*, edited by Brad Hooker and Margret Olivia Little, 48–78. Oxford: Oxford University Press.
Raz, Joseph. 2001. *Value, Respect, and Attachment*. Cambridge: Cambridge University Press.
Regan, Donald H. 2004. "Why Am I My Brother's Keeper?" In *Reason and Value: Themes from the Philosophy of Joseph Raz*, edited by R. Jay Wallace, Samuel Scheffler, and Michael Smith, 202–30. Oxford: Clarendon Press.
Ridge, Michael. 2011. "Reasons for Action: Agent-Neutral vs. Agent-Relative." The Stanford Encyclopedia of Philosophy *Winter*: http://plato.stanford.edu/archives/win2011/entries/reasons-agent.
Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.
Rønnow-Rasmussen, Toni. 2011. *Personal Value*. Oxford: Oxford University Press.
Rosati, Connie. 2008. "Objectivism and Relational Good." *Social Philosophy and Policy* 25: 314–49.
Rowland, Richard. 2013. "Wrong Kinds of Reasons and Consequences." *Utilitas* 25: 405–16.
Rowland, Richard. 2016. "In Defence of *Good Simpliciter*." *Philosophical Studies* 173: 1371–95.
Scanlon, Thomas. 1998. "*What We Owe to Each Other*." Cambridge, MA: Belknap.
Scheffler, Samuel. 1982. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Clarendon Press.
Scheffler, Samuel. 2004. "Projects, Relationships, and Reasons." In *Reason and Value: Themes from the Philosophy of Joseph Raz*, edited by R. Jay Wallace, Samuel Scheffler, and Michael Smith, 247–69. Oxford: Clarendon Press.
Schroeder, Mark. 2007a. "Teleology, Agent-Relative Value, and 'Good.'" *Ethics* 117: 265–95.
Schroeder, Mark. 2007b. *Slaves of the Passions*. Oxford: Oxford University Press.
Schroeder, Mark. 2009. "Buck-Passer's Negative Thesis." *Philosophical Explorations* 12: 341–47.
Seglow, Jonathan. 2013. *Defending Associative Duties*. New York: Routledge.
Smith, Michael. 2003. "Neutral and Relative Value after Moore." *Ethics* 113: 576–98.
Smith, Michael. 2009. "Two Kinds of Consequentialism." *Philosophical Issues* 19: 257–72.
Suikkanen, Jussi. 2009. "Buck-Passing Accounts of Value." *Philosophy Compass* 4/5: 768–79.
Tan, Kok-Chor. 2010. "Equality and Special Concern." *Canadian Journal of Philosophy* 40: 73–98.
van Willigenburg, Theo. 2005. "Reason and Love: A Non-Reductive Analysis of the Normativity of Agent-Relative Reasons." *Ethical Theory and Moral Practice* 8: 45–62.
Väyrynen, Pekka. 2006. "Resisting the Buck-Passing Account of Value." *Oxford Studies in Metaethics* 1: 295–324.
Velleman, David. 1999. "Love as a Moral Emotion." *Ethics* 109: 338–74.
Wallace, R. Jay. 2010. "Reasons, Values, and Agent-Relativity." *Dialectica* 64: 503–28.

Cite this article: Löschke, J. 2020. Agent-Relative Reasons as Second-Order Value Responses. *Canadian Journal of Philosophy* 50: 477–491, doi:10.1017/can.2020.3