

Normative Reasons Qua Facts and the Agent-Neutral/Relative Dichotomy: a Response to Rønnow-Rasmussen

Jamie Buckland¹

Received: 6 April 2016 / Revised: 21 July 2016 / Accepted: 12 September 2016 /

Published online: 4 October 2016

© Springer Science+Business Media Dordrecht 2016

Abstract This paper offers a defence of the distinction between agent-neutral and agent-relative reasons for action from scepticism aired by Toni Rønnow-Rasmussen. In response it is argued that the Nagelian notion of an agent-neutral reason is not incomprehensible, and that agent-neutral reasons can indeed be understood as obtaining states of affairs that count in favour of anyone and everyone performing the action they favour. Furthermore, I argue that a distinction drawn between agent-neutral and agent-relative reason-statements that express the salient features of reason-constitutive states of affairs is neither reductive in the sense of reducing normative reasons to the propositional content of an agent's mental state, nor trivial in the sense of locating the distinction merely in an agent's description of the world.

Keywords Agent-neutral · Agent-relative · Normative reasons · Motivating reasons · Explanatory reasons · Thomas Nagel · Toni Rønnow-Rasmussen

Since its conception by Thomas Nagel the distinction between agent-neutral and agent-relative reasons for action has played a dominant role in normative theorizing. In a nutshell, agent-neutral reasons are said to be reasons for anyone and everyone. For instance, the *general* reasons there are for us all to help those in need or to protect the environment. Agent-relative reasons, on the other hand, are said to be reasons only for particular individuals. For instance, the *special* reasons there are for each of us to look after our own interests, or the interests of our family and friends, etc. (Nagel 1970: Ch. XI, 1986: Ch. IX).

As far as philosophical distinctions go, this is remarkably straightforward. Nevertheless, the ability to draw this distinction successfully has recently been questioned by Toni Rønnow-Rasmussen (2009, 2011 Ch. 9, 2012). Working within a framework which takes

✉ Jamie Buckland
jamie.buckland.philosophy@gmail.com

¹ The Nova Institute of Philosophy, New University of Lisbon (IFILNOVA), Faculdade de Ciências Sociais e Humanas, Universidade Nova de Lisboa, Av. de Berna, 26 - 4º Piso, 1069-061 Lisbon, Portugal

normative reasons to be constituted by facts understood as obtaining states of affairs, he raises two concerns. Firstly, that no sense can be made of the notion of an agent-neutral reason; the obtaining states of affairs which constitute agent-neutral reasons are not in any way about the agent for whom they are reasons, so their existence is “incomprehensible” (Rønnow-Rasmussen 2011: 141). And, secondly, that the leading method of understanding the distinction in terms of agent-neutral or agent-relative reason-statements that express reason-constitutive states of affairs is trivial; such a distinction concerns itself only with what agents *take to be* their normative reasons qua the propositional content of their mental states, not the normative reasons themselves. Consequently, the distinction is located merely in how agents describe the world rather than delineating anything of normative or evaluative significance (Rønnow-Rasmussen 2012: 100). This paper amounts to a substantial response to this scepticism, and a defence of a non-trivial distinction between agent-neutral and agent-relative reasons for action qua obtaining states of affairs.

1 Reasons Qua Facts and the Essentialist Sense of the Distinction

In order to understand the claim that agent-neutral reasons are incomprehensible we need to get clear on two things. Firstly, that Rønnow-Rasmussen takes normative reasons for action to be constituted by “facts favouring action or the adoption of an attitude”¹ (Rønnow-Rasmussen 2011: xi). And, secondly, that his scepticism concerning a successful distinction between agent-neutral and agent-relative reasons is directed specifically towards Nagel’s “*essentialist sense*” of the distinction understood in terms of “what does and what does not *essentially* refer to a particular person [as] the *owner of the reason*” (Rønnow-Rasmussen 2011: 129, 2012: 97 (emphasis added)):

If a reason can be given a general form which does not include an essential reference to the person who has it, it is an *agent-neutral* reason. [...] If on the other hand the general form of a reason does include an *essential reference* to the person who has it, it is an *agent-relative* reason (Nagel 1986: 152–153).

From here, Rønnow-Rasmussen captures the essentialist account of the distinction in the following way (where R represents the statement expressing the fact that constitutes the reason; x the agent for whom the fact that R expresses a reason; and φ the requirement of the reason, i.e., what the fact that R expresses counts in favour of x doing, promoting, desiring, intending, or feeling):

If R states a reason *for* x *to* φ , then:

R states an *agent-relative* reason for x *if and only if*.

R contains essential reference to x ,

Otherwise, R states an agent-neutral reason for x to φ .

¹ The precise nature of these facts will be examined below. At this stage it’s sufficient to note that Rønnow-Rasmussen takes facts to be “synonymous” with ontologically substantial or “*thick*” obtaining states of affairs rather than ontologically “*thin*” propositional entities (Rønnow-Rasmussen 2011: 135–136, 2012: 98–99).

Moreover,

R contains an essential reference to x if and only if.

R is not logically equivalent with any other statement Q that does not refer to x (Rønnow-Rasmussen 2011: 130).

For example:

(AN) “A person y is drowning” (in circumstance C)

(AR) “My daughter y is drowning” (in circumstance C)

These reason-statements express an obtaining state of affairs which may figure in a *that-clause* representing answers to questions of the form: “Why should I/you φ ” (e.g. jump in the water)? (Rønnow-Rasmussen 2012: 97; Alvarez 2010: 41–42) “A person y is drowning” is said to express an agent-neutral reason for someone to help y , whereas “My daughter y is drowning” is said to express a first-personal agent-relative reason for y ’s parents to help y – a *different* or *special* reason from the agent-neutral reason an agent may have who is not y ’s parent² (Rønnow-Rasmussen 2011: 130, 2012: 97).

2 Against Agent-Neutral Reasons: the Personalizability Implication Feature

The above captures the essentialist sense of the distinction in terms of the reasons-qua-facts idea. Nevertheless, Rønnow-Rasmussen insists we should be sceptical of this distinction by virtue of what he coins the “*personalizability implication feature*” of normative reasons:

Since all reasons are apparently for someone to φ , and a reason to φ is only a reason for someone if it somehow refers to this someone, it follows that all reasons to φ are in their very form reasons that refer to the person who has the reason to φ . This, in its turn is just another way of saying that all reasons to φ are, on entirely formal grounds, agent-relative reasons (Rønnow-Rasmussen 2011: 132).

In light of this personalizability implication feature, the claim is that agent-neutral reasons (understood as obtaining states of affairs) are incomprehensible:

On the essentialist approach, given the truth of an agent-neutral reason-statement, there is an obtaining state of affairs such that it is a reason for x . Moreover, it is a reason for x despite the fact that the state of affairs is not in any way about x . [...]

² Talking in terms of a *different* reason can be slippery; there’s a perfectly good sense in which the fact that y is drowning constitutes the same reason for anyone: a reason for anyone who can to help y . Indeed, the ultimate crux of my argument is that the distinction between agent-neutral and agent-relative reasons does not amount to a distinction between two kinds of fact, but rather two ways a reason (understood as a fact) can vary in its scope of application.

The question “Why is this fact a reason for x ?” cannot therefore be answered by pointing to some feature of the situation that even in a minimal sense concerns x . This is quite remarkable. The conclusion we appear to be obliged to draw is that there is nothing about the fact that makes it a reason for x , but the fact is nonetheless a reason for x (Rønnow-Rasmussen 2011: 137 (emphasis added)).

It will emerge in due course why the personalizability implication feature does not render the notion of an agent-neutral reason incomprehensible. Though in order to see this we first need to understand why Rønnow-Rasmussen is wrong to direct this scepticism explicitly towards Nagel’s notion of an agent-neutral reason for action.

3 Nagelian Agent-Neutrality

In understanding Nagel’s notion of an agent-neutral reason the first thing to note is that Nagel does not share Rønnow-Rasmussen’s commitment to what has become known as *reasons primitivism*: the idea that the concept of a normative reason defies non-circular definition, i.e., that reasons count in favour of acts or attitudes by providing reasons for them (Rønnow-Rasmussen 2011: xi; Scanlon 1998: 97; Parfit 2011: 31). Nagel could certainly allow that reasons supervene on obtaining states of affairs, but for Nagel reasons are formulated as predicates derived from universal practical principles and *defined* in the following way:

Every reason is a predicate R , such that for all persons x and all events φ , if R is true of φ , then x has prima facie reason to promote φ (Nagel 1970: 47).

For every token reason there is a corresponding predicate R such that it figures in a universally quantified normative proposition providing prima facie reason for agents to promote the occurrence of particular events ranging from specific actions, inactions, circumstances, to more general states of affairs, outcomes, or ends. This constitutes a definition of a reason understood as a count noun (a non-normative fact or non-normative predicate) in terms of *reason* understood as a mass noun: that which explains the reason’s prima facie normative attraction (Broome 2013: 63). In Nagel’s principle-based sense then, normative reasons are understood as both an explanation and a justification of action, i.e., a normative reason for x to φ is what explains why x ought to φ (Nagel 1970: 14–15).

The formal distinction between agent-neutral and agent-relative reasons emerges within the universally bound reason-predicate, R :

Formally, [an agent-relative reason] is one whose defining predicate R contains a *free occurrence* of the variable x . (The free agent-variable will, of course, be free only *within* R ; it will be bound by the universal quantification over all persons which governs the entire formula.) All universal reasons and principles expressible in terms of the basic formula either contain a free agent-variable or they do not. The former are [agent-relative]; the latter will be called [agent-neutral]³ (Nagel 1970: 90).

³ Nagel originally drew this distinction in terms of “objective” and “subjective” reasons for action, though later adopted Derek Parfit’s agent-neutral/relative terminology (see Parfit 1984: 27, 143).

Taken alongside the essentialist sense of the distinction found in *The View from Nowhere*, it's clear that it should be understood as a distinction between the general form of those reasons which contain an essential reference to the agent who has the reason and those reasons which lack such reference.

Reconsidering the above examples can bring out the differences between the two distinctions. Suppose I find myself in the unfortunate situation in which my daughter is drowning and conclude there's a reason for me to help her, and suppose I express this reason via the first-personal agent-relative reason-statement "My daughter y is drowning." As Rønnow-Rasmussen understands the distinction, "for a fact to be an agent-relative reason [...] it would be for some obtaining state of affairs to involve the agent for whom the fact is a reason." Provided this statement is true, i.e., the state of affairs does obtain, then the obtaining state of affairs expressed "is in some sense about the agent for whom it operates as a reason" (Rønnow-Rasmussen 2011: 137). In the above circumstances then, the feature of the situation which makes this fact into a reason for me is *that it is my daughter who is in need of help*. This is the salient feature of the obtaining state of affairs which counts in favour of me helping my daughter by providing me with a reason to do so.

Nagel's notion of a reason for action is more complex than this. Rather than accepting the idea that reasons defy a non-circular analysis, he seeks to address the question: "What reason do I have to help my daughter?" where my initial answer may be something like: *that the act will prolong my daughter's life*. However, given that Nagel considers it a "bear condition" that reasons are universal, the reason there is for me to help my daughter must also be expressible in general terms, which may (among other things) amount to the following responses⁴:

- (1) That the act will prolong the life of my child;
- (2) That the act will prolong the life of Jamie Buckland's child;
- (3) That the act will prolong the life of someone's child.

All of these can be understood as general reasons valid for all persons, but only (1) is an agent-relative reason because the anaphoric pronoun "my" is functioning as a free agent-variable. This takes the following general form alongside its corresponding universal principle:

- (1) (x, φ) (If φ will prolong the life of x 's child, then x has reason to promote φ .)

Everyone has reason to do what will prolong the life of *his or her own* children.

For all persons x and events φ , if the predicate "... will prolong the life of x 's child" is true, then this constitutes an agent-relative reason only for x to φ : "the free agent-variable prevents the transmission of derivative influence to acts of any other person" (Nagel 1970: 93). Consequently, agent-relative reasons have a restricted scope of

⁴ It should be stressed here that universality found in the general form of Nagel's reasons is merely logical and imposes no rational constraint on choices of action. It amounts merely to the idea that, as long as we take ourselves to have reasons, then a judgement about these reasons entails claims about the reasons other agents would have under relevantly similar circumstances (see Mackie 1977: 83–102; Scanlon 1998: 74–75).

application; they are reasons only for particular individuals to promote the events to which they apply.

Agent-neutral reasons are different. The reasons expressed in (2) and (3) contain no free occurrence of the agent-variable x within the antecedent of the reason-predicate:

(2) (x, φ) (If φ will prolong the life of Jamie Buckland's child, then x has reason to promote φ .)

Everyone has reason to do what will prolong the life of Jamie Buckland's child.⁵

(3) (x, φ) (If $(\exists z)$ (φ will prolong the life of z 's child), then x has reason to promote φ .)

Everyone has reason to do what will prolong the life of *someone's* child.

In these instances the agent-variable does not occur within the reason-predicate, rendering any reference to the agent for whom the consideration is a reason closed, determined, and therefore *inessential*. The key idea is that there is no open reference to the doer of the act within the antecedent of the reason-predicate. Reference to an agent occurs within the requirement of the reason, i.e., the event which the reason counts in favour of the agent promoting, but the reference is closed; it is not a particular agent-denoting reference because the reason was not particularized to anyone in its antecedent. Consequently, there is (putatively) an agent-neutral reason for *anyone* to promote the occurrence of the event in question.⁶ The crucial observation to make, then, is that unlike agent-relative reasons, agent-neutral reasons are *shared*; they give all agents a common aim, goal, or end⁷ (Parfit 1984: 27). Indeed, given that Nagel's distinction fundamentally concerns the manner in which reasons can vary in their scope of application, the crucial exegetical point to make against Rønnow-Rasmussen's interpretation of Nagel's distinction is that Nagel's distinction is *not* "silent about the *range* of agents for whom something is a reason" (Rønnow-Rasmussen 2011: 129). On the contrary, an essential reference to the agent for whom the consideration is a reason is tantamount to restricting the range of agents to which the reason applies.⁸

Now that we understand how Nagel's distinction is supposed to work, we can ask ourselves whether we can draw a similar dichotomy if we understand reasons as

⁵ It might be objected here that if I, Jamie Buckland, claim that my reason for φ -ing is that φ will prolong Jamie Buckland's child's life, then it is – by definition – an agent-relative reason. After all, there is reference to Jamie Buckland in both the consequent of the conditional and its antecedent, i.e., all we have is (1) but x is substituted for Jamie Buckland. However, its full representation does not contain a *free* agent-variable, so the reason remains agent-neutral. This includes me (Jamie Buckland) qua member of anyone and everyone, but the point is that from an agent-neutral standpoint I needn't know that I am, in fact, Jamie Buckland.

⁶ Those who are not in a position to *do* anything about the matter should not interfere with the occurrence of the event in question or, at the very least, desire that the event occur.

⁷ There is of course a *sense* in which agent-relative reasons are shared reasons, for at one level of description the agent-relative reason each of us has to prolong the lives of our own children is a reason for anyone to do the same thing, i.e., what prolongs the lives of our children. But since what is best for one agent's child might not be best for another's, then the reasons we have may well be opposed.

⁸ Nagel and Parfit's accounts of the distinction should *not*, then, be understood as distinct; they are extensionally equivalent (see Ridge 2011: Section 1; cf. Rønnow-Rasmussen 2011: 131–132, 2012: 69).

obtaining states of affairs.⁹ For Rønnow-Rasmussen the obtaining state of affairs in which my daughter y is drowning is said to constitute an agent-relative reason for me to help her. Likewise, the obtaining state of affairs in which his daughter is drowning is said to constitute an agent-relative reason for him to help his daughter. Universally speaking then, for all agents x , the obtaining state of affairs in which their daughter is drowning is an agent-relative reason for them to help their daughter:

(x, y) (The obtaining state of affairs in which) x 's daughter y is drowning is a reason for x to help y .

The reason is agent-relative in virtue of the free occurrence of the agent-variable x within the antecedent of the reason-predicate. Consequently, the reason has a restricted scope of application. Likewise, for all agents x , the obtaining state of affairs in which a person y is drowning is said to constitute an agent-neutral reason for anyone (who can) to help y :

(x, y) (The obtaining state of affairs in which) a person y is drowning is a reason for x to help y .

Again, there is an implicit reference to an agent here, but only within the requirement of the reason. In this sense there is nothing “mystical” or “incomprehensible” about agent-neutral reasons, they’re simply reasons for anyone who happens to fall within their wide scope of application.

In stating that agent-neutral reasons are not reasons for specific individuals it may look as if I am denying the personalizability implication feature. Indeed, Rønnow-Rasmussen describes this kind of response as a far-fetched refusal to take up the challenge it poses (Rønnow-Rasmussen 2011: 139). The obtaining states of affairs that constitute reasons, he insists, cannot be “merely” reasons, or “reasons period”, because they do not simply “call for” action, but call for action by agents in a position to perform that action. However, not specifying a particular agent for whom the obtaining state of affairs is a reason does not render the notion of an agent-neutral reason “mystical”, for characterizing all normative reasons as obtaining states of affairs in the first place implies that they could, at least in principle, constitute reasons for *someone* to act – *what else could a normative, practical reason for action be for?* For instance, say that it’s impossible for anyone to help my daughter in circumstance C . Does this entail there is *no reason* for her to be helped, i.e.,

⁹ Aside from a casual endorsement of Joseph Raz’s observation that “reasons have a vague and incomplete criteria of identity” (Raz 2006: 109). Rønnow-Rasmussen offers no direct support in favour of the idea that reasons are obtaining states of affairs rather than true propositions. Support comes from two indirect worries. Firstly, propositions and their corresponding reason-statements express merely what agents *take to be* (*believe to be*) their reasons, rather than the reasons themselves. And, secondly, that propositions and the reason-statements that express them are only the “tip of the iceberg” as far as *complete* normative reasons are concerned: the idea that “there is one and only one thing I have a reason to do at t^1 and that [...] is determined by how the world precisely is at t^1 ” (Rønnow-Rasmussen 2011: 146). We’ll consider these ideas in more detail below, but it’s worth stressing at this point that I have no problem with the idea that obtaining states of affairs can constitute or ground reasons for action, *per se*. As I’ll argue below, the point is that unlike Rønnow-Rasmussen I do not regard this idea to be in contention with the propositional account. On another note, the claim that *all* normative reasons for action are obtaining states of affairs or natural worldly facts just looks false. For instance, if asked why I helped someone, I may say that they wouldn’t have been able to achieve some goal or that they may have suffered unnecessarily unless I helped them. Such counterfactuals seem to provide reasons for action, but are not obtaining states of affairs. This is not, however, a criticism I’ll explore here.

that no reason-constitutive entity exists? Not on the view that normative reasons are constituted by worldly obtaining states of affairs. It simply entails that no one was in a position to help my daughter, which is why we might describe the event as a tragedy, rather than saying there was no reason for her to be helped in the first place. By no means does this amount to the “admission that reasons for action can, quite literally, be nobody’s reason.” Rather, there was a reason for anyone to help my daughter – the obtaining state of affairs in which she was drowning – but unfortunately no one was there to respond to it.

At this stage of my argument I conclude that there are no grounds for Rønnow-Rasmussen’s scepticism concerning the comprehensibility of agent-neutral reasons for action. The personalizability feature implication poses no threat to the comprehension of agent-neutral reasons understood as obtaining states of affairs or otherwise.

4 A Slippery Distinction?

With agent-neutral reasons successfully demystified, it’s time to turn our attention to Rønnow-Rasmussen’s attack on the notion of an agent-relative reason. To describe this as an attack might sound odd in light of his claim “that, given all plausible assumptions, all normative reasons are agent-relative” (Rønnow-Rasmussen 2011: xi, 91). However, I have argued that the personalizability implication feature points only towards a trivial sense of relativity captured within the requirement of the reason. In this sense scepticism regarding the possibility of drawing a non-trivial and philosophically interesting distinction between agent-neutral and agent-relative reasons still stands. Indeed, the task now is to respond to further arguments which state that a distinction drawn between agent-neutral and agent-relative *reason-statements* which express reason-constitutive obtaining states of affairs is located trivially in the way we describe the world rather than distinguishing anything of normative or evaluative significance (Rønnow-Rasmussen 2012: 105–106; also see Ridge 2011: Section 2).

Recall the pair of reason-statements from above:

(AN) “A person *y* is drowning” (in circumstance *C*)

(AR) “My daughter *y* is drowning” (in circumstance *C*)

Rønnow-Rasmussen found the agent-neutrality or agent-relativity of these reason-statements confusing because their status appears to depend on the agent to whom they addressed. For instance, if the statement “A person *y* is drowning” is addressed to *y* it seems to express an agent-relative reason for *y*. Likewise, if the statement “My daughter *y* is drowning” is addressed to an agent, *z*, then it seems to express an agent-neutral reason for *z* to help *y*. In light of this observation, Rønnow-Rasmussen draws the following conclusion:

[G]iven that we have in mind *a certain kind of reason* [...] it is clear that whether a reason-statement expresses something agent-relative or agent-neutral depends, in part, on the person to whom it is addressed. On the other hand, if the truth of a reason-statement is not relativized to the way the person to whom it is addressed understands the statement, the above confusion concerns something else, namely what people *take to be* reasons (Rønnow-Rasmussen 2011: 130–131 (emphasis added)).

Without further qualification it's tricky to grasp what this worry amounts to. Indeed, the final sentence is particularly difficult to parse. In order to get a grip on this alleged slipperiness we need to get clear on the distinction Rønnow-Rasmussen is drawing between the *certain kinds of reasons* – the truth of whose corresponding neutral/relative reason-statements depends on the agent to whom they are addressed – and what agents *take to be* their reasons – the truth of whose corresponding agent-neutral or agent-relative reason-statement is not relativized to the way the person to whom it is addressed to understands the statement.

5 Decomposing Reasons: Normative, Motivating, and Explanatory

The *certain kind of reasons* Rønnow-Rasmussen has in mind are what we've been referring to as normative reasons: obtaining states of affairs that count in favour of doing, believing, or desiring something by providing a reason for them (Rønnow-Rasmussen 2011: xi). Let's call these *objective normative reasons* (objective in the sense that the reason exists independently of the subjective belief states of the agents for whom they are reasons):

Objective Normative Reason: If an obtaining state of affairs R stands in a favouring-relation to φ , then R is a normative reason for x to φ .

For instance, a reason for x to jump in the river was the fact that his daughter was drowning. For Rønnow-Rasmussen, these normative or “*real*” reasons must be distinguished from what an agent *takes to be* their reasons, their “*apparent*”, “*operative*”, or “*motivating*” reasons understood as the “owned” propositional content of an agent's belief state: “a state of affairs which as a matter of fact is *believed* to be a (normative) reason by the agent.”¹⁰

¹⁰ This notion of “*owned*” propositional content is crucial to understanding Rønnow-Rasmussen's account of motivating reasons; the idea being that propositions are always *someone's* because they are what a certain person might express at a given time and place. What is notable, however, is that Rønnow-Rasmussen does not distinguish clearly between *propositional entities* and *propositional attitudes* (the mental state held by an agent towards a proposition). At one point he cites the importance the intentional context of a proposition has on an agent's behaviour in terms of a propositional attitude, but it is often ambiguous as to whether he is referring to reason-statements expressing the propositional attitude, i.e., *my belief that my daughter y is drowning*, or the “owned” proposition \langle that *my daughter y is drowning* \rangle qua some abstract entity (Rønnow-Rasmussen 2012: 100). As we'll see now, Rønnow-Rasmussen makes a particular effort to distinguish motivating reasons qua the propositional content of an agent's belief state from explanatory reasons qua the agent's belief state itself, insisting that it's an agent's *belief that R* (in combination with an appropriate desire) which causes them to act, not some “proposition-like entity” (Rønnow-Rasmussen 2011: 134). However, he goes on to maintain that motivating and explanatory reasons are both reductive in the same sense. For instance, in relation to reducing normative reasons to “in-the-head propositional entities” (motivating reasons) he insists: “What is in the head is more accurately (or, at any rate, just as correctly) described as *the grasping of the fact or feature, but not the fact or feature itself*.” Yet, he also speaks of reducing normative reasons to what goes on in the agent's mind, in terms of propositional attitudes (a mental state held towards a proposition) (Rønnow-Rasmussen 2012: 106–107). The problem with this ambiguity is that the distinction Rønnow-Rasmussen draws between motivating reasons qua the propositional content of an agent's mental state and explanatory reasons qua the mental state itself collapses. He insists the propositional content of an agent's mental state expresses an obtaining state of affairs an agent believes to be a (normative) reason for them to φ . Yet simultaneously holds that propositional content does not explain action, rather, it is *my belief that my daughter y is drowning* which explains why I act. However, if motivating reasons are just *what agents believe* to be their reasons qua grasped in-the-head propositional entities, then this serves equally well in explaining action. I'll expand and defend this idea in Section 7.

(Rønnow-Rasmussen 2011: 133–134 (emphasis added); Persson 2005: Ch. 8). Let's call these *subjective motivating reasons* (subjective in the sense of the reason existing in virtue of the propositional content of the agent's belief state):

Subjective Motivating Reason: If x believes that R stands in a favouring relation to φ , then R is a motivating reason for x to φ .

For instance, x had a reason to jump into the water because he believed that his daughter was drowning.

In relation to the above concerns, for Rønnow-Rasmussen this means that the reason-constitutive element of a reason-statement such as “My daughter y is drowning” or “A person y is drowning” is not an ontologically “thin” propositional entity such as \langle that my daughter y is drowning \rangle or \langle that a person y is drowning \rangle qua “the content of possible judgements that are true” (Rønnow-Rasmussen 2012: 99). Rather, the reason is constituted by an ontologically concrete entity consisting of worldly properties and objects:

The state of affairs is concrete [...] in the sense that it contains a number of features that may be abstracted from these entities. Some features are more salient than others and what, e.g., AN and AR express is these salient features rather than the whole picture. On the thick approach the propositional element in a reason-statement is like the tip of the iceberg; it displays only a tiny fraction of the whole entity (Rønnow-Rasmussen 2012: 98).

Given this substantial account of the facts that constitute normative reasons then, the normative reason *there is* for me to help my daughter is understood as an entity with extrinsic features which relate to that which makes the proposition \langle that my daughter y is drowning \rangle true – its metaphysical *truth-maker*.

We'll consider the distinction between normative reasons qua obtaining states of affairs and motivating reasons qua propositional content of thoughts that are true in more detail in Section 7, but it's worth noting at this point that Rønnow-Rasmussen's position amounts to a form of *nonreductive realism* about normative reasons, detaching them from Davidsonian talk concerning the psychological belief/desire states of the agent which cause an agent to act (Davidson 1963). Speaking somewhat crudely, let's call these *explanatory reasons* (i.e., explanation in causal terms of):

Explanatory Reason: If x believes that R , and this belief in conjunction with x 's desires would normally cause x to φ , then x 's belief that R is an explanatory reason why x might φ .¹¹

For instance, if I believe that my daughter y is drowning and desire to help her, then (akrasia aside) the combination of *my belief that my daughter y is drowning* and my desire to help her is the causal explanation of why I act (should I do so): “it's the agent's belief that R , rather than *the fact that R* alone, that explains why she acted in the relevant way” (Rønnow-Rasmussen 2011: 133). So, although it's a combination of an

¹¹ I thank an anonymous reviewer for this improved version of my original formulation.

agent's mental belief/desire states which *causes* them to act (not the propositional content of their mental state), it's Rønnow-Rasmussen's resistance towards closing the gap between normative and explanatory reasons that draws him to the objective account of normative reasons. Normative reasons for action are obtaining states of affairs that get their truth conditions independently of an agent's psychological states, hence the normative and the explanatory are treated as separate vectors. Consequently, an agent's motivating reasons must not be confused with the mental belief/desire states which cause them to act, nor with the obtaining states of affairs that constitute the normative reasons for them to act. For Rønnow-Rasmussen, motivating reasons are not explanatory reasons, and neither are they normative (Rønnow-Rasmussen 2011: 134).

My own view on this matter is that normative questions concerning the justification of action do not lose their significance when we attend to questions concerning what agents believe to be their (normative) reasons. Indeed, I'll offer some defence of this idea in Section 7. Nevertheless, now that we're clear on how the notion of a practical reason is said to decompose into normative, motivating, and explanatory varieties, we can return to the confusion regarding the owners and addressees of reason-statements discussed in Section 4.

6 Getting a Grip: the Essentially Indexical Account of Agent-Relative Reasons

There appear to be two concerns. The first is that statements expressing reason-constitutive states of affairs do not express truths about real normative reasons; they merely express the propositional content of an agent's mental state – an obtaining state of affairs the agent takes to be a (normative) reason. Therefore, the issue concerning whether a statement expresses an agent-relative or an agent-neutral reason depending on the agent to whom it is addressed is an issue concerning an agent's apparent or motivating reasons, not normative or real reasons (Rønnow-Rasmussen 2011: 134, n. 13). The second concern is that an account of the distinction that understands reasons in terms of the propositional content of an agent's mental states locates the difference between agent-neutral and agent-relative reasons in our descriptions of the world, rather than delineating anything of normative or evaluative significance, and is, therefore, trivial (Rønnow-Rasmussen 2012: 105). Let's dispense with the first worry now, for the natural response to my resolution involves addressing the triviality objection in Section 8.

Recall Rønnow-Rasmussen found it strange that the truth of a statement expressing a reason-constitutive obtaining state of affairs can be relativized to how the person to whom it is addressed to understands the statement. In terms of the now qualified propositional account of motivating reasons, Rønnow-Rasmussen supposes that if I, qua owner of the proposition *<that my daughter y is drowning>* were to address or direct this proposition to another agent, *z*, in the form of the expressed statement "My daughter *y* is drowning", then this allegedly agent-relative reason should be regarded as an agent-neutral reason for someone other than myself to help my daughter. Likewise, if another agent, *z*, qua owner of the proposition *<that a person y is drowning>* were to address this to *y* in the form of the expressed statement "A person *y* is drowning", then this allegedly agent-neutral reason should be regarded as an agent-relative reason for *y* to get help.

In responding to this concern, the first thing to note is that there's nothing peculiar about the situation envisaged above. I can easily imagine myself standing on the bank of a river, witnessing my daughter drowning, formulating the proposition $\langle \textit{that my daughter } y \textit{ is drowning} \rangle$ and screaming: "My daughter y is drowning!" at bystander z . And it's equally conceivable that, upon hearing this statement and witnessing my daughter drowning, z might jump into the river to help her. What is odd, however, is to suggest that one might think that *the entity* that is the reason *for* z to jump into the river and help my daughter is the owned proposition $\langle \textit{that my daughter } y \textit{ is drowning} \rangle$ qua the content of *my* judgment that is true (i.e., *my* motivating reason). Rather, we can grant that the entity which constitutes the reason for us both to help y is the ontologically concrete obtaining state of affairs in which y is drowning. The relevant difference, however, is that z does not stand in the same relationship to the reason-constitutive obtaining state of affairs as I do, i.e., it is not *his* daughter who is drowning. It is this relational difference the distinction between agent-neutral and agent-relative reasons is trying to capture.

For Rønnow-Rasmussen, reasons for action are ontologically concrete obtaining states of affairs that count in favour of acts or attitudes by providing reasons for them. However, where some ontologically concrete state of affairs does obtain, an agent must still recognise that it does in fact obtain in order to be able to consider whether it constitutes a reason for himself or anyone else to do anything. And, when an agent recognises that some feature of an obtaining state of affairs is an agent-relative reason for them to act, we can follow John Skorupski in saying that the salient feature of the obtaining state of affairs is presented to them in an agent-relative or an *essentially indexical* way (Skorupski 2010: 63; Pettit 1988: 165). For instance, although the state of affairs in which my daughter is drowning *just is* the state of affairs in which *someone's* daughter is drowning, or in which Jamie Buckland's daughter is drowning, in recognising this fact as a reason *for me* to act I recognise it indexically (provided I know that I am someone whose daughter is drowning). It is this indexical feature of the situation, as I find it, which enables me to express my reason agent-relatively via a declarative reason-statement containing an anaphoric pronominal back-reference to myself – the agent for whom the state of affairs is a reason. Indeed, Rønnow-Rasmussen's own argument against the comprehensibility of agent-neutral reasons depended on our ability to cite such salient features in ordinary language via the expression of true reason-statements.¹²

By Rønnow-Rasmussen's own admission, the reason-statement "My daughter y is drowning" expresses a salient feature of the concrete obtaining state of affairs in which my daughter is drowning. It is this "propositional feature" – $\langle \textit{that it is my daughter who is in need of help} \rangle$ – which "makes" the obtaining state of affairs into an agent-relative reason *for me*. When the same state of affairs constitutes a reason for z to help y , the relevant features of the situation are not presented to z in this indexical way.¹³ The salient feature of the state of affairs in which y is drowning that make this into an agent-neutral reason for z is simply $\langle \textit{that a person } y \textit{ is drowning} \rangle$. Granted, if z were asked

¹² Recall Section 2 (Rønnow-Rasmussen 2011: 137).

¹³ It must be stressed here that we are concerned simply with the status of normative reasons qua obtaining states of affairs. In the sense of generating action both agent-neutral and agent-relative reasons are, in John Perry's terms, "essentially indexical" because agents will have to believe that they (reflexively) ought to act at time t . Nevertheless, the reason-constitutive fact *itself* needn't be essentially indexical (Perry 1979).

why he jumped into the river to help *y*, he might express the proposition “*His* daughter was drowning” making a pronominal back-reference to me. Likewise, if *z* happens to know my daughter is drowning, he could also direct the statement expressing the proposition “Your daughter is drowning” to me, but neither of these statements contains a pronominal back-reference to *himself* qua agent for whom the consideration is a reason. The point is that although the reason-statement is made true and effectuated by the salient features identified by the propositional content of a judgement that is true, this need not entail that the propositional element itself constitutes the reason for action.¹⁴ Rather, the idea is that one-and-the-same obtaining state of affairs can constitute a reason which can vary in its scope of application depending on an agent’s relation to it. Agent-neutral reasons are reasons for anyone and everyone. Their constitutive states of affairs count in favour of anyone who can perform that which the state of affairs counts in favour of doing because they appeal to general features of the situation. Agent-relative reasons, on the other hand, are special reasons only for those individuals who stand in a special relationship to the constitutive states of affairs identifiable by the use of either first or third-personal pronominal back-reference to the agent for whom they are reasons.

7 Recomposing Reasons

At this point Rønnow-Rasmussen will likely press the following concern:

If we are suspicious about the idea that a statement expressing an alleged reason will in itself settle the question whether there are any agent-neutral/relative reasons, then consideration of the proposition expressed by this statement will hardly offer further assistance. And if the true proposition does not settle the question – “does the statement express a normative reason?” – I do not see how the aspect, i.e., the truth-maker of the proposition, can do so. Something more is needed (Rønnow-Rasmussen 2011: 136).

It is, however, frightfully unclear what, exactly, “*something more*” could be referring to here. One suggestion relates to Rønnow-Rasmussen’s idea that reasons understood as the propositional content of judgements that are true express only the “tip of the iceberg” as far as “*complete*” normative reasons are concerned: the notion that “there is one and only one thing I have a reason to do at t^1 and that this is determined by how the world precisely is at t^1 ” (Rønnow-Rasmussen 2011: 146). Nevertheless, though conceptually plausible, the idea of a complete reason is both epistemically and practically useless – it’s impossible to know what one’s own or anyone else’s complete reason could be; consequently, it’s impossible to act for such a reason, i.e., it’s impossible for such a reason to constitute a practical reason for action. As Rønnow-Rasmussen himself notes: “the sort of entity that it suggests is constitutive of a reason would simply involve features that it would be impossible to include in the

¹⁴ It could be argued that this understanding of propositions is “Russellian”, i.e. that propositions themselves are instantiations of concrete properties. So, the proposition <*that my daughter y is drowning*> consists of my daughter *y*, drowning, and the “being-in” relation between them (Russell 1903: 47; also see Suikkanen 2012: 598).

supervenience base of a reason” (Rønnow-Rasmussen 2012: 110). The obtaining state of affairs that constitutes a reason for an agent to help y at t^1 cannot be inclusive of the fact \langle that a person y is drowning \rangle along with how the world precisely is at t^1 for this would entail that one’s reason to help y would somehow be inclusive of the fact \langle that a cat is currently eating a smoked salmon dinner in Canada \rangle , or the fact \langle that a supernova is occurring within the Eagle Nebula \rangle some 7000 light-years from earth.

This objection seems decisive. Indeed, the very idea that normative, practical reasons for action could be both epistemically and practically inaccessible is an odd one to say the least.¹⁵ Nevertheless, in response to this “*Too Inclusive Objection*” Rønnow-Rasmussen has suggested that the supervenience base of a reason will also include a fact about which facts are relevant or irrelevant:

What makes something into a complete reason are certain salient features, together with the peculiar feature that any remaining feature is irrelevant. We do not have to give up the idea that among the reason making characteristics of the world, we should only include those that are relevant for the explanation of why the agent ought to ϕ . It is just that one of these relevant features is the feature that there are no other relevant features. So if at a given time the facts are such that certain features positively call for some action ϕ by me, and it is the case that there are no other features that would disregard this call, then I have a complete normative reason for acting (Rønnow-Rasmussen 2012: 112).

Yet, rather than lending support to the idea that one’s normative reasons must be constituted by how the world precisely is at t^1 , this method for accounting for complete normative reasons just collapses into the propositional account we’ve been discussing, i.e., certain features of an ontologically concrete obtaining state of affairs, identifiable by propositions, count in favour of some act or attitude ϕ by someone. What’s more, provided these propositions are true, they’re no less “normative” than their metaphysical truth-makers. This idea can be brought out via further consideration of the connection between so-called normative and motivating reasons.

Any comprehensive study of the distinction between agent-neutral and agent-relative reasons must address the question: “What is a reason for action?” And, as we’ve already seen, in answering this question Rønnow-Rasmussen cites two “varieties” of reason: objective normative reasons qua concrete obtaining states of affairs (the *real* reasons *there are* for agents to act – independently of what they believe), and subjective motivating reasons qua the propositional content of thoughts that are true (the reasons agents *have* which depend on what they believe). The upshot of this nonreductive realism was that motivating reasons were considered neither normative nor explanatory: when I jump into the river and help my daughter the normative reason *for which* I act is the obtaining state of affairs in which my daughter is drowning, and it is *my belief that my daughter y is drowning* (combined with my desire to help her) which explains why I jump into the river and help her, not the proposition \langle that my daughter y is drowning \rangle qua the content of my mental state. However, it doesn’t seem to be the case that

¹⁵ At one point Rønnow-Rasmussen goes as far as suspecting that “reasons need not be graspable” (Rønnow-Rasmussen 2012: 111). Though again, if a normative reason for action need not be graspable by anyone then I no longer understand how it can be said to constitute a practical reason for anyone to do anything.

justificatory questions concerning normative reasons decompose when we attend to questions concerning what agents believe they have a reason to do. Indeed, the decomposition of reasons into the varieties of normative, motivating, and explanatory reasons just looks false, for there is a perfectly good sense in which motivating reasons qua propositional content figure in “reason-*why*” explanations and, when true, inherit their normativity from their constitutive states of affairs (Hornsby 2008).

Say I’ve accidentally ingested a number of hallucinogenic mushrooms and undergo an experience of my daughter drowning, i.e., I have a false belief with the propositional content \langle *that my daughter y is drowning* \rangle . In this instance I take this to be a normative reason for me to help my daughter, i.e., the proposition expresses a state of affairs I *believe* to be a (normative) reason for me to jump into the river. Nevertheless, for Rønnow-Rasmussen this cannot be the reason *for which* I act; in knowing what an agent believes we merely know the *reason why* they act as they do, and, should I believe \langle *that my daughter y is drowning* \rangle , then according to Rønnow-Rasmussen, *that I believe that my daughter y is drowning* explains why I act as I do. However, this is more than just a causal explanation. The reason why the milk went sour was because it was left out of the refrigerator, but the milk didn’t *have* a reason to go sour in the explanatory sense that I *had* a reason to jump into the river. Indeed, when you witness me ingesting magic mushrooms and hours later jumping into the river screaming: “My daughter *y* is drowning!” you treat me as someone who is acting for what they *take to be* a (normative) reason. Indeed, it’s the normativity of this reason that explains why unlike the milk going sour it can be the subject of rational assessment. If I jump into the river for the reason *that I believe that my daughter y is drowning*, then from my first-personal perspective there is a (normative) reason for me to do so: \langle *that my daughter y is drowning* \rangle . Moreover, if this proposition were in fact true, i.e., it expressed the salient feature of the obtaining state of affairs in which my daughter was drowning, then the proposition, \langle *that my daughter y is drowning* \rangle , inherits its normativity from its constitutive state of affairs. By no means does this amount to the admission that “just because we believe something to be a reason, it *is* a reason”, entail a commitment to “a crude sort of relativism”, or reduce reasons to “what goes on in an agent’s head in terms of propositional attitudes” (Rønnow-Rasmussen 2011: 134, 2012: 106–107). The point is simply that when I believe the proposition \langle *that my daughter y is drowning* \rangle is, for me, a reason to act, there is, from my perspective, a reason for me to jump into the river which *would* have been the reason *for which* I acted had she actually been drowning.

8 Is the Propositional Account Trivial?

Unlikely as Rønnow-Rasmussen is to accept the recomposed account of reasons I’ve just offered, in relation to the propositional account of the distinction I defended in Section 6, he will press the concern that such a distinction is *trivial*. The worry being that a distinction drawn between agent-neutral and agent-relative reason-statements that express true propositions “locates the difference between agent-neutral and agent-relative reasons only in our descriptions of the world and does not, therefore, express anything of normative or evaluative significance” (Rønnow-Rasmussen 2011: 136, 2012: 105).

In an attempt to elucidate this concern more formally, Rønnow-Rasmussen has appealed to the following schema:

$$(4) x [R] z (\varphi)$$

Unlike the three-place characterization of the essentialist account we considered in Section 1, this schema is comprised of four elements, where focus shifts from reason-statements that express reason-constitutive states of affairs to reason-constitutive propositions that encompass the essential “owner element”: R represents the reason-constitutive proposition qua the content of a thought that is true (its “reason-maker”); x denotes the agent entertaining the proposition (its “owner”); z refers to the agent for whom the proposition is a reason; and φ is the act/attitude which z has a reason to perform/hold in light of the true proposition R entertained by x . With this set, Rønnow-Rasmussen maintains the agent-relative reason for him (TRR) to help his drowning dog takes the following form:

$$(5) TRR [My dog is drowning] TRR (\text{saving of the dog by } TRR)$$

The reason is said to be agent-relative because “*the proposition essentially refers to the person who issued one of these statements, and thus endorsed the proposition*”, thereby containing an ineliminable cross-reference between the person referred to as “my” and the person for whom the proposition is a reason (Rønnow-Rasmussen 2012: 100 (emphasis added); 103). However, we should stress here that it needn’t be the case that x and z must refer to the same individual for a reason-constitutive proposition to be considered agent-relative. Take, for instance:

$$(6) x [\text{Your dog is drowning}] TRR (\text{saving of the dog by } TRR)$$

Here x is entertaining a proposition she takes to be an agent-relative reason for TRR to save his dog. What’s important in this third-personal case then, is that x ’s proposition $\langle \text{that your dog is drowning} \rangle$ contains an ineliminable back-reference to the agent for whom it is a reason, and *not* to the “owner” of the proposition.¹⁶ Nevertheless, in order for TRR to grasp this proposition as a reason for him to save his dog, he must grasp that it is *his* dog that is drowning. Agent-neutral reasons, by contrast, are said to lack a cross-reference between the reason-constitutive proposition and the agent for whom the proposition is a reason. For instance, the proposition owned by x $\langle \text{that a dog is drowning} \rangle$ is said to constitute an agent-neutral reason for TRR to φ (Rønnow-Rasmussen 2012: 103–104):

$$(7) x [\text{A dog is drowning}] TRR (\text{saving of the dog by } TRR)$$

Likewise, x ’s proposition $\langle \text{that } TRR\text{'s dog is drowning} \rangle$ is also said to constitute an agent-neutral reason for TRR to φ :

$$(8) x [TRR\text{'s dog is drowning}] TRR (\text{saving of the dog by } TRR)$$

¹⁶ I thank the anonymous reviewer for pressing me on this point.

With this schematization in place, Rønnow-Rasmussen raises the following objection.¹⁷ Consider a claim expressing the proposition “Buster is drowning”:

(9) x [Buster is drowning] TRR (saving of Buster by TRR)

This proposition is said to be an agent-neutral reason-maker for TRR to save Buster. However, if (9) is true, then (5) must also be true, i.e., x 's proposition “Buster is drowning” and TRR 's proposition “My dog is drowning” refer to one-and-the-same concrete obtaining state of affairs in which a dog owned by TRR is drowning, this, in turn, trivializes the distinction:

[T]he thin account has to say, as a minimum, that there is a normative difference between [(5), (7), and (9).] It should normatively matter, in the sense that it should at least be feasible that only one of these might be a reason for TRR , or at least be a stronger/weaker reason for TRR than the alternatives. Otherwise the distinction would not be interesting or important. However, I fail to see that such a normative difference is detectable, and so I find the distinction trivialized (Rønnow-Rasmussen 2012: 105–106).

But the idea that it should be feasible that only one of these might be a reason for TRR , or at least be a stronger or weaker reason for TRR than the alternatives is simply wrongheaded, for this is not something a formal dichotomy between agent-relative and agent-neutral reasons could ever tell us.¹⁸ Indeed, “normative” and “evaluative” questions concerning, say, the weight of reasons can only be settled once we have a distinction that's useful to work with and (presumably) a substantive theory of what is valuable, i.e., a normative theory about how agents should respond to the obtaining states of affairs that constitute their reasons given their relation to them, what is required or expected of them, and their underlying value systems, etc.

Of course the distinction between agent-neutral and agent-relative reasons concerns itself with normative matters, but this is precisely what the use of reason-statements expressing the salient features of reason-constitutive states of affairs is designed to illuminate. It highlights a distinction between how reasons can be said to vary in their scope of application. Whether or not you think the distinction is normatively significant in the grand scheme of things is a separate question, but this is something that needs to be argued for independently. You may, for instance, ask yourself as William Godwin once did: “What magic is there in the pronoun ‘my’ that it should justify us in overturning the decisions of impartial truth,” i.e., what *difference* (if any) does it make

¹⁷ Notice that the proposed schema is asymmetrical in its treatment of agent-relative and agent-neutral reasons. Given the propositional account of reasons under scrutiny, reasons are understood as the owned propositional content of an agent's mental state – an obtaining state of affairs an agent believes to be a (normative) reason for *them* to ϕ . However, unlike the agent-relative reason expressed in (5), the agent-neutral reasons expressed by (7) and (8) are not representative of the first-personal motivating reasons we've been considering, i.e., (7) and (8) do not express the propositional content of TRR 's judgements that are true. Rather, they express propositions owned by x which x believes to be reasons for TRR to ϕ . And, as we've seen in (6), it is possible to capture agent-relativity in this schema via the use of third-personal pronouns. Qua reasons *for* x to ϕ , however, (6), (7), (8), and (9 (below)) are all agent-neutral.

¹⁸ Nagel was always keen to stress the formal nature of the original distinction, in particular the idea that the formal condition of agent-neutrality could not dispense with a substantive theory of value (Nagel 1970: 126).

that it is *my* child, or *your* child, or *his* child who is drowning in circumstance *C*, rather than an individual of no significant relation to us? And you may conclude, as Godwin once did, that it makes no difference at all, or you may conclude, as Bernard Williams once did, that it makes *all* the difference (Godwin 1798; Williams 1981; MacIntyre 1983; also see deGaynesford 2010). Nevertheless, substantive normative questions concerning the magic or indeed the weight of essentially indexed agent-relative reasons can only be asked (let alone settled) once we've identified them. And, as the history of ethics has shown, this is by no means a trivial observation.

9 Conclusion

This paper has demonstrated that Rønnow-Rasmussen's scepticism concerning a successful distinction between agent-neutral and agent-relative reasons understood as obtaining states of affairs is ungrounded. I have argued that agent-neutral reasons qua obtaining states of affairs can indeed be understood as reasons for anyone and everyone, and that we can draw a non-trivial distinction between agent-neutral and agent-relative reasons in terms of statements that express the salient features of these obtaining states of affairs identifiable with (though not reducible to) the propositional content of an agent's belief state.¹⁹ Rather than showing the distinction between agent-neutral and agent-relative reasons to be located in trivial descriptions of the world which lack normative or evaluative significance, Rønnow-Rasmussen has failed to recognize that the formulation of reason-statements via linguistic or logical means is necessary to capture the simple intuition that certain states of affairs can be said to express special reasons in the sense of being relativized to the agent for whom they are reasons. Our linguistic ability to express one and the same reason-constitutive obtaining states of affairs in both indexical and non-indexical terms does not render the distinction between agent-relative and agent-neutral reasons trivial or unhelpful – even if we subscribe to an ontological conception of reasons which maintains there are no distinctively indexical facts or irreducibly indexical propositions or states of affairs (cf. Ridge 2011: Section 2). The indexical might not be an objective feature of the obtaining states of affairs which constitute one's reasons for action. Nevertheless, it's often a feature of reasons as an agent finds them, and it is reasons as agents find them which both motivate and explain action.

Acknowledgements An earlier version of this paper was presented to the Laboratory of Ethics and Political Philosophy within the Nova Institute of Philosophy, Universidade Nova de Lisboa. I am grateful to all attendees for their comments and suggestions. I would also like to thank the two anonymous reviewers for *Philosophia* for their extremely helpful comments and criticisms that improved the paper immensely. Additionally, thanks are owed to Emma Borg, Max de Gaynesford, Brad Hooker, Toni Rønnow-Rasmussen, Philip Stratton-Lake, and Alan Thomas for their helpful comments and discussion on earlier work from which this paper was derived.

¹⁹ Another method of defending the notion of an essentially indexical agent-relative reason is to argue that reasons-statements are non-extensional contexts (Suikkanen 2012). However, since I have defended a propositional-based account of reasons, and propositions are themselves non-extensional, Suikkanen's observations complement my own position.

References

- Alvarez, M. (2010). *Kinds of reason*. Oxford: Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. Chichester: Wiley-Blackwell.
- deGaynesford, R. M. (2010). The bishop, the valet, the wife, and the ass. In J. Cottingham & B. Feltham (Eds.), *Partiality and impartiality in ethics*. Oxford: Oxford University Press.
- Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23), 685–700.
- Godwin, W. (1798). *Enquiring concerning political justice and its influence on modern moral and happiness* (3rd ed.) (Eds.) Kraminick, I. (1985) London: Penguin.
- Hornsby, J. (2008). A disjunctive conception of acting for reasons. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: perception, action, knowledge*. Oxford: Oxford University Press.
- MacIntyre, A. (1983). The magic in the pronoun "My". *Ethics*, 94(1), 113–125.
- Mackie, J. L. (1977). *Ethics: inventing right and wrong*. London: Penguin Books.
- Nagel, T. (1970). *The possibility of altruism*. Oxford: Oxford University Press.
- Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On what matters: volume one*. Oxford: Oxford University Press.
- Perrson, I. (2005). *The retreat of reason: a dilemma in the philosophy of life*. Oxford: Oxford University Press.
- Perry, J. (1979). The problem of the essential indexical. *Noûs*, 13(1), 3–21.
- Pettit, P. (1988). The paradox of loyalty. *American Philosophical Quarterly*, 25(2), 163–171.
- Raz, J. (2006) The trouble with particularism (Dancy's version). *Mind*, Vol. 115, 99–120.
- Ridge, M. (2011) Reasons for action: agent-neutral vs. agent-relative. *The Stanford Encyclopedia of Philosophy*. Zalta, E. N. (ed.), URL = <http://plato.stanford.edu/archives/n2011/entries/reasons-agent/>.
- Rønnow-Rasmussen, T. (2009). Normative reasons and the agent-neutral/relative dichotomy. *Philosophia*, 37(2), 227–243.
- Rønnow-Rasmussen, T. (2011). *Personal value*. Oxford: Oxford University Press.
- Rønnow-Rasmussen, T. (2012). Reasons and two kinds of fact. *Discusiones Filosóficas*, 13(20), 95–113.
- Russell, B. (1903). *The principles of mathematics*. New York: Norton.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge, MAS: The Belknap Press of Harvard University.
- Skorupski, J. (2010). *The domain of reasons*. Oxford: Oxford University Press.
- Suikkanen, J. (2012). Reason-statements as non-extensional contexts. *The Philosophical Quarterly*, 62(248), 592–613.
- Williams, B. (1981) *Persons, character, and morality*. In *Moral Luck*. Cambridge: Cambridge University Press.