

Skorupski and Broome on the Agent-Neutral/Relative Distinction

JAMIE BUCKLAND

University of York

I have two aims in this article. The first is to break the deadlocked exchange between John Skorupski and John Broome concerning how best to understand Thomas Nagel's distinction between agent-neutral and agent-relative reasons for action. The second is to provide a reformulation of the distinction which captures an uncontroversial distinction between those reason-giving considerations which encapsulate an indexical relationship between an agent and an object of moral concern, and those which do not. The resolution of this exchange, and subsequent reformulation of the dichotomy, has two important ramifications for contemporary debates in moral philosophy. First and foremost, it reveals the true, pre-theoretical nature of the distinction between agent-neutral and agent-relative reasons for action, and how the notion of agent-relativity cannot be utilized to underwrite the existence of deontic constraints. And, second, it provides definitive support for Skorupski's claim that agent-relative reasons are not the defining feature of deontological ethics.

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. However, the role the distinction has played has often been one of confusion, rather than clarity; superficial similarity of terminology disguises the fact that philosophers use the terms in radically different ways due to their competing aims and prior assumptions. This article examines one instance of this confusion, and offers a resolution to the deadlocked exchange between John Skorupski and John Broome concerning how best to understand Nagel's formulation of the distinction. After breaking this deadlock and discussing a number of problems that beset the Nagelian version of the distinction, I'll offer a reformulation of the dichotomy which captures a pre-theoretical distinction between agent-neutral and agent-relative reason-giving facts or principles. I'll conclude the article by citing some implications this view has for contemporary moral theorizing.

I. NAGEL'S PRINCIPLE-BASED DISTINCTION

To understand how Nagel's distinction works, we need to get clear on how he understands the concept of a normative (or justifying) reason for action.

For Nagel, all normative reasons can be formulated as predicates derived from universal practical principles, and are defined in the following way:

[E]very 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 .¹

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 events, ranging from specific actions, inactions, circumstances, to more general states of affairs, outcomes, or ends. In this sense, a reason can apply to a specific event such as my drinking a beer ten seconds from now, or it can apply to a more general state of affairs, such as someone being in good health. In the latter case, any number of more specific events may promote the end in question.²

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

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

¹ *The Possibility of Altruism* (Oxford, 1971), pp. 47–8. Nagel uses ‘ p ’ for the agent variable, and ‘ A ’ for the act, event, or circumstance to be promoted. I have amended all schemas to remain consistent throughout this article.

² Notice that this constitutes a definition of a *normative 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, *Rationality Through Reasoning* (Oxford, 2013), p. 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 promote φ is what explains why x ought to promote φ (*Possibility*, pp. 14–15). This differs from the now popular view known as *non-reductive realism*, which takes the notion of a normative reason to be primitive, and explicitly circular in definition. On the non-reductive view, a normative reason for action is simply a consideration (fact) which counts in favour of some act or attitude by providing a reason for it (see T. M. Scanlon, *What We Owe to Each Other* (Cambridge, 1998), p. 71; Derek Parfit, *On What Matters*, vol. 1 (Oxford, 2011), p. 31; Toni Rønnow-Rasmussen, *Personal Value* (Oxford, 2011), p. xi; cf. Jamie Buckland, ‘Normative Reasons qua Facts and the Agent-Neutral/Relative Dichotomy’, *Philosophia* 45 (2017), pp. 207–25). Nagel’s position is more complex than this; he envisages an intimate connection between the circumstances (or facts) that provide agents with reasons to act, and the way in which action is explained by reasons brought under the control of practical normative principles. For the moment, I will work explicitly with Nagel’s concept of a normative reason. I will, however, explain how this relates to the non-reductive view in section VII.

³ Nagel, *Possibility*, p. 90. Nagel originally drew the distinction in terms of ‘objective’ and ‘subjective’ reasons for action, though he later adopted Parfit’s agent-neutral/

With this set, we can utilize Nagel's distinction to delineate a formal difference between the egoist's and the utilitarian's fundamental principles of action:

Egoist: Everyone has reason to do what has a propensity to benefit themselves.

(x, φ) (If φ has a propensity to benefit x , then x has reason to promote φ .)

Utilitarian: Everyone has reason to do what has a propensity to maximize aggregate utility.

(x, φ) (If φ has a propensity to maximize aggregate utility, then x has reason to promote φ .)

As far as these examples go, the significant mark of agent-neutrality is whether or not *the reason*, i.e. the fact picked out by the predicate (not the act or event to be promoted), can be specified without an essential reference to the agent for whom it is a reason. So, for any given egoist, if, say, eating healthily has a propensity to benefit him, then that gives him an agent-relative reason to promote its occurrence (i.e. to eat healthily). The reason is agent-relative in the sense that its defining predicate contains an essential reference to the agent for whom it is a reason (i.e. it is the agent *himself* who stands to benefit from eating healthily). On the other hand, for any given utilitarian, if donating money to charity has a propensity to maximize aggregate utility, then that gives her an agent-neutral reason to promote its occurrence (i.e. to donate money to charity). The reason is agent-neutral in the sense that its defining predicate lacks reference to the agent for whom it is a reason (i.e. it is simply that aggregate utility will be maximized by performing the act in question, not, specifically, *the agent's* utility).

II. AGENT-NEUTRALITY: PERFORMING ACTIONS, OR PROMOTING EVENTS?

Now that we're clear on how Nagel's distinction is supposed to work, we can begin to mediate the exchange that occurred between John Skorupski and John Broome concerning how best to understand the notion of agent-neutrality.

Skorupski has always insisted he is using the terms 'agent-neutral' and 'agent-relative' to 'coincide respectively' with how Nagel uses the

agent-relative terminology (see Parfit, *Reasons and Persons* (Oxford, 1984); Nagel, *The View from Nowhere* (Oxford, 1986), p. 125).

terms ‘objective’ and ‘subjective’ in *The Possibility of Altruism*.⁴ There is, however, a subtle difference between the two distinctions which, for Broome, has substantial implications for how we understand the notion of an agent-neutral reason for action.

To see this, consider Skorupski’s most recent statement of the distinction:

(It’s being the case) that $R\varphi$ gives x reason to φ .

Where ‘ x ’ is replaceable by terms denoting agents and ‘ φ ’ by terms denoting actions open to x . (That is, for any given x , ‘ φ ’ ranges over the action-types open to x .) If all instances of the schema are true, then ‘ R ’ is a *reason-giving predicate on agents and actions open to the agent*.⁵ Now we define agent-neutral and agent-relative reasons as follows:

- (i) If ‘ R ’ contains a free occurrence of ‘ x ’ then it is an agent-relative predicate. If it does not, it is an agent-neutral predicate.
- (ii) A reason for action which is expressible by an agent-neutral predicate is agent-neutral. A reason for action which is not so expressible is agent-relative.⁶

Now consider the following reason for performing an action:

That the act has a propensity to benefit an agent who does it.

For Broome, the predicate ‘... has a propensity to benefit an agent who does φ ’ expresses an egoistic reason for an agent to directly perform an act which has a propensity to benefit himself – a reason that is

⁴ Skorupski, *The Domain of Reasons* (Oxford, 2011), p. 63.

⁵ It’s important to stress that, for Skorupski, φ ranges over *action-types* open to x , and that types of action can be open to more than one agent, i.e. ‘in more than one agent’s choice-situation’ (Skorupski, ‘Neutral versus Relative: A Reply to Broome, and McNaughton and Rawling’, *Utilitas* 8 (1996), pp. 235–48, at 236).

⁶ Skorupski, *Domain*, pp. 63–4. One objection to this approach is that it renders all reasons agent-relative reasons ‘for the simple fact that any reason for an agent to perform an act must appeal to some fact about the agent’s performing that act’ (Douglas W. Portmore, ‘Agent-Relative vs. Agent-Neutral’, *International Encyclopedia of Ethics*, ed. Hugh LaFollette (2013), p. 3). For example, Douglas Portmore has insisted that the fact that Jane is in need of some good cheer is an agent-neutral reason for him to throw her a surprise party. Yet, when this reason is fully specified in Nagelian terms, it emerges, incorrectly, as agent-relative:

(x) (It’s being the case) that x ’s throwing Jane a surprise party would ensure that Jane experiences some good cheer gives x a reason to throw Jane a surprise party.

However, the fact that any reason for an agent to perform an act must appeal to some fact about the agent’s performing that act is a trivial feature of *all* reasons for action, and is not, therefore, an interesting mark of ‘agent-relativity’ (see Michael Ridge, ‘Reasons for Action: Neutral vs. Relative’, *The Stanford Encyclopedia of Philosophy*, <<https://plato.stanford.edu/archives/fall2017/entries/reasons-agent/>> (2017). Indeed, as I stressed at the end of [section I](#), the significant mark of agent-relativity is whether or not the reason-defining predicate (not the act or event to be promoted) contains an essential reference to the agent for whom it is a reason.

uncontroversially agent-relative.⁷ However, when this predicate is run through Skorupski's machinery, it emerges as agent-neutral:

- (1) (x, φ) (It's being the case) that doing φ has a propensity to benefit an agent who does φ gives x reason to φ .

This predicate contains no free occurrence of x , so it is agent-neutral according to Skorupski's definition. Yet it expresses an egoistic reason for action, and egoistic reasons should definitely not be counted as agent-neutral. Skorupski's definition is too broad, then.⁸

Broome then insists that Nagel's version of the distinction – formulated in terms of *promoting the occurrence of an event* – bypasses this problem:

- (2) (x, φ) (It's being the case) that doing φ has a propensity to benefit an agent who does φ gives x reason to promote φ .

Understood in this way, agent-neutral reasons require agents to share a common *aim*.⁹ For instance, say you and I are at a party, and there is one beer left in the ice bucket. And, suppose that anyone at the party has reason to drink it (for the reason that drinking it has a propensity to benefit an agent who does so). According to Broome, Skorupski must define this reason as *both* agent-neutral and egoistic because it applies to anyone who can perform it:

- (3) (x) (It's being the case) that drinking the last beer has a propensity to benefit an agent who drinks it gives x reason to drink the last beer.

Yet, (3) gives each agent at the party the conflicting aim of drinking the last beer themselves, and cannot, therefore, be *genuinely* agent-neutral. Formulating the distinction in terms of a reason to promote the occurrence of the event in which an agent drinks the last beer bypasses this problem:

- (4) (x) (It's being the case) that drinking the last beer has a propensity to benefit an agent who drinks it gives x reason to promote the occurrence of an agent drinking the last beer.

This (putatively) gives *anyone* a reason to promote the occurrence of the event in which an agent drinks the last beer, not merely

⁷ Broome, 'Skorupski on Agent-Neutrality', *Utilitas* 7 (1995), pp. 315–17, at 315. Broome's actual predicate reads: 'that the act gives an expected benefit to an agent who does φ '. Again, I have amended the predicate to remain consistent.

⁸ Broome, 'Skorupski', p. 315.

⁹ Parfit, *Reasons*, p. 104.

an agent who has a propensity to benefit from performing the act herself. Hence, Broome's insistence that 'agent-neutrality cannot be satisfactorily defined in terms of reasons to do an act'.¹⁰

III. MEDIATING THE DISPUTE

In response, Skorupski maintained that Broome had provided a defective analysis of his own formulation of the distinction by treating it asymmetrically to Nagel's, i.e. Broome's (4) is not, in fact, an accurate representation of Nagel's position. Rather, the genuine parallel, in Nagelian terms, would actually take the following agent-neutral form:

- (5) (x, φ) (It's being the case) that *promoting* φ has a propensity to benefit an agent who promotes φ gives x reason to *promote* φ .

Adding:

A problem I have with (5) is that I don't know what it is to promote an action . . . And, now whatever talk about 'promoting' may mean – the point Broome makes about (1) could be made about (5). The predicate 'promoting φ has a propensity to benefit an agent who promotes φ ' contains no free occurrence of x and so it is agent-neutral according to Nagel's position. Yet it expresses an egoistic reason for action, and egoistic reasons should definitely not be counted as agent-neutral. Nagel's definition is too broad then.¹¹

Now, because Skroupski fails to understand what it means to promote an action, this leaves only two possible readings of Broome's predicate ' . . . has a propensity to benefit an agent who does φ ' – both of which are agent-neutral, and neither of which expresses an egoistic reason for action. On the first reading, 'an agent' could refer to some (existentially quantified) agent or other whose doing φ has a propensity to benefit himself:

- (6) $(x, \varphi) (\exists y)$ (It's being the case) that y 's doing φ has a propensity to benefit y gives x reason to φ .

Or, on a second reading, 'an agent' could refer to any (universally quantified) agent whose doing φ has a propensity to benefit himself:

- (7) (x, y, φ) (It's being the case) that y 's doing φ has a propensity to benefit y gives x reason to φ .

In (6), we have an agent-neutral reason for any given agent to directly perform an action-type because (for the reason that) performing that action-type (the action-type in question) has a propensity to benefit

¹⁰ Broome, 'Skorupski', p. 317.

¹¹ Skorupski, 'Neutral Versus', pp. 237–8.

some agent or other who performs it. For instance, if y 's drinking the last beer has a propensity to benefit y , then that (putatively) gives anyone reason to drink the last beer (for the reason that it benefits y). This resembles the way in which an annoying sibling might eat the last slice of cake simply to upset their younger sister, even though they dislike the cake themselves – a reason that is spiteful, but not egoistic. And, in (7), we have an agent-neutral reason for any given agent to perform the action-type in question because performing that action-type has a propensity to benefit anyone who performs it. For instance, if drinking the last beer has a propensity to benefit anyone who drinks it, then that gives anyone reason to drink it (for the reason that it has a propensity to benefit anyone who drinks it). Nevertheless, although the egoist will accept the truth of (7), as Skorupski points out, it does not express an egoistic reason for action; the egoist will accept (7), but only because:

- (8) (y) (y 's doing φ has a propensity to benefit y) & doing φ is open to $x \rightarrow x$'s doing φ has a propensity to benefit x .

What brings out his adherence to egoism is that he has to establish that doing φ gives an expected benefit to x before he considers x to have reason to do φ . The point about rational egoists is not that they reject all agent-neutral principles; the point is that they agree with such principles only in so far as they follow from the agent-relative principle which they endorse.¹²

But what if we think in terms of reasons for an agent to *promote* φ ? At this point, it's vital to remember that Nagel does not speak in terms of reasons to promote acts. Rather, he speaks in terms of reasons to promote the occurrence of *events* (ranging from specific events (actions) to more general states of affairs). With this in mind, consider, again, Broome's predicate:

That the act has a propensity to benefit an agent who does φ .

As we've just seen, this reason-defining predicate utilizes an ambiguous quantifier in its reference to '*an agent who does φ* '. This is problematic for Broome, because it's entirely a matter of interpretation how ambiguous reason-defining predicates are to be understood – that's precisely Nagel's point.¹³ Broome takes the predicate to express an agent-relative reason for an agent to directly perform an act which has a propensity to benefit himself. But, if this is the case, then its formulation will have to contain a free occurrence of the agent-variable

¹² Skorupski, 'Neutral Versus', p. 237.

¹³ Nagel, *Possibility*, p. 90.

within the antecedent of its defining predicate (where, in the limiting case, if φ is an act, doing φ counts as promoting φ):

- (9) (x, φ) (It's being the case) that doing φ has a propensity to benefit x gives x reason to promote [i.e. *do*] φ .

However, Nagel's point is that the same predicate can also be subjected to an *agent-neutral interpretation*, which would, in this case, express a reason for *anyone* to promote the occurrence of the event (act) that gives an expected benefit to an agent who does it – *whether or not that agent is oneself*:

- (10) (x, φ) (It's being the case) that doing φ has a propensity to benefit an agent who does φ gives x reason to promote an agent doing φ .

In this instance, the predicate lacks an occurrence of the agent-variable x : $R(\varphi)$ ' φ has a propensity to benefit an agent'. Or, more specifically: $(\exists y)$ (y 's doing φ has a propensity to benefit y):

- (11) (x, φ) (It's being the case) $(\exists y)$ that y 's doing φ has a propensity to benefit y gives x reason to promote y 's doing φ .

This predicate contains only the bound agent-variable, y , and still applies when x and y are the same person.¹⁴ For instance, if y 's drinking the last beer has a propensity to benefit y , then that gives *anyone* reason to promote the occurrence of y drinking the last beer. Granted, in their primary application, there will always be significant overlap between agent-neutral and agent-relative reasons. But, for Nagel, the derivative influence of the agent-neutral reason extends to acts which promote the occurrence of acts by *other* agents which have a propensity to benefit *them*:

[Agent-neutral reasons] are not just universal reasons in the sense that anyone can have them; they are in addition reasons for anyone to promote what they apply to. They are not reasons *for* particular individuals, but simply for the *occurrence* of things they hold true. . . . The reasons then apply in the first instance to conduct which meets these descriptions, but they also apply, derivatively, to actions which promote more conduct of the same kind, whether in oneself or in others.¹⁵

¹⁴ Nagel is explicitly clear that he has this kind of existential quantification in mind (*Possibility*, p. 92).

¹⁵ Nagel, *Possibility*, pp. 91, 93.

Hence, Nagel's substantive claim that mere acknowledgement of the idea that all agent-relative reasons have agent-neutral counterparts is enough to explain the *possibility* of altruism.¹⁶

IV. THE NECESSARY ASYMMETRY

It's now clear that Skorupski's use of the terms 'agent-neutral' and 'agent-relative' does not 'coincide respectively' with how Nagel uses the terms 'objective' and 'subjective' in *The Possibility of Altruism*. Nevertheless, taken alongside Broome's concerns, this does not undermine the formal adequacy of Skorupski's version of the distinction. On the contrary, it enables us to observe an asymmetry in Nagel's formulation of the distinction via his differential treatment of those agent-neutral reasons which are generated when we subject our agent-relative reasons to the requirement of agent-neutrality (the agent-neutral counterparts) and those reasons which are fundamentally agent-neutral in content.

As I explained in note 16, the substantive claim of *The Possibility of Altruism* was that all agent-relative reasons had to be subsumable under their agent-neutral counterparts. For instance, if you accept that you have an agent-relative reason to perform an act which has a propensity to benefit yourself, then you must be able to understand this reason as not, essentially, your own. Additionally, you must also recognize an agent-neutral reason for anyone (including yourself) to

¹⁶ Nagel, *Possibility*, pp. vii–viii, pp. 96–7. The original claim was that *all* agent-relative reasons had to be *subsumed* under their agent-neutral counterparts, although Nagel eventually abandoned this idea in light of objections raised by Nicholas Sturgeon ('Altruism, Solipsism, and the Objectivity of Reasons', *Philosophical Review* 83 (1974), pp. 374–402) and Stephen Darwall ('Nagel's Argument for Altruism', *Philosophical Studies* 25 (1974), pp. 125–30; *Impartial Reason* (Ithaca, 1983)). The idea was that the agent who acknowledges only agent-relative reasons is committed to a practical version of solipsism called *dissociation*: an 'inability to make practical judgements about other persons in the same sense in which one can make them in one's own case' (*Possibility*, p. 107). Unless another agent's reasons have *motivational content* for yourself from an impersonal (agent-neutral) standpoint, then you cannot recognize them as a practical agent in the same sense that you recognize yourself as a practical agent. Moreover, because you're simply a mere agent amongst others, you would also be unable to recognize your own reality from that impersonal standpoint. Understood in this way, the *requirement of agent-neutrality* is functioning as a formalized version of Rawls's veil of ignorance, accompanied by a Sidgwickian conception of objectivity (John Rawls, *A Theory of Justice*, rev edn. (Harvard, 1999), pp. 118–22; Henry Sidgwick, *The Methods of Ethics* (Indianapolis, 1874/1981), pp. xx, 82, 420; Christopher Tollefsen, 'Sidgwickian Objectivity and Ordinary Morality', *The Journal of Value Enquiry* 33 (1999), pp. 57–70; Alan Thomas, *Thomas Nagel* (Acumen, 2009), pp. 126–8). Again, if you take yourself to have agent-relative reasons to *do* anything at all, then you must be able to understand these reasons as not, essentially, your own. From the 'Point of View of the Universe', you must also recognize an agent-neutral reason for anyone (including yourself) to *promote* what the reason applies to (*Possibility*, pp. 16–17).

promote what the reason applies to (recall the agent-relative (9) and its agent-neutral counterpart (11)):

- (9) (x, φ) (It's being the case) that doing φ has a propensity to benefit x gives x reason to promote [i.e. do] φ .
 (11) (x, φ) (It's being the case) $(\exists y)$ that y 's doing φ has a propensity to benefit y gives x reason to promote y 's doing φ .

Now, although Nagel characterizes all reasons in terms of reasons to promote the occurrence of events, ultimately, all reasons direct suitable agents to perform action-types which promote the event in question.¹⁷ As we saw in section I, the egoist has an agent-relative reason to *do* whatever has a propensity to benefit himself, and the utilitarian has an agent-neutral reason to *do* whatever has a propensity to maximize aggregate utility. At this level, then, agent-neutral reasons are equally concerned with agents performing action-types which promote the occurrence of a particular state of affairs. However, one cannot characterize the agent-neutral counterpart reasons generated by agent-relative reasons in the same way without generating reasons that are in direct tension with the original notion of agent-neutrality, i.e. a universally shared reason for anyone to promote the occurrence of the event in question. Recall (6), for instance:

- (6) $(x, \varphi) (\exists y)$ (It's being the case) that y 's doing φ has a propensity to benefit y gives x reason to φ .

Drawn in terms of reason to perform the action-type in question, (6) gives any given agent reason to perform the action-type that y has reason to do (for the reason that it benefits y), which can result in conflicting aims. Contra Broome's conclusion, then, it is *not* that 'agent-neutrality cannot be satisfactorily defined in terms of reasons to do an act'. The point is subtler: *the agent-neutral counterpart reasons which are generated when we subject our agent-relative reasons to the Sidgwickian objectivity constraint (the requirement of agent-neutrality) cannot be satisfactorily defined in terms of reasons to perform the action-type in question.* Indeed, if Nagel did not characterize agent-neutral counterparts in terms of promoting the occurrence of the event in question, then the central argument of *The Possibility of Altruism* would be untenable – the counterpart reasons would not reflect the possibility of altruistic action. There is, then, an asymmetry in Nagel's original formulation of the distinction. Like agent-relative reasons,

¹⁷ This includes *doing* nothing. If an agent is not in a position to do anything to promote the occurrence of the event in question, then they should, at least, *desire* that the event occur (*Possibility*, p. 91).

those agent-neutral reasons which are fundamentally agent-neutral in content *can* be successfully defined in terms of reasons for a given agent to perform the action-type in question. However, the agent-neutral counterpart reasons which are generated by subjecting our agent-relative reasons to a Sidgwickian objectivity constraint cannot be satisfactorily defined in terms of performing the action-type in question – they *must* be defined in terms of promoting the occurrence of the event (action-type) in question.¹⁸

At this stage of my argument, I conclude that Skorupski's formulation of the distinction is more viable than Nagel's, namely in its ability to offer a symmetrical distinction between agent-neutral and agent-relative reasons drawn in terms of reasons for agents to perform action-types that are open to them.¹⁹ Nevertheless, there are still some ambiguities surrounding the Nagelian version of the distinction that need to be ironed out, particularly in relation to the status of deontic constraints.

V. THE STATUS OF DEONTIC CONSTRAINTS

Stemming from Nagel's infamous discussion of the issue in *The View from Nowhere*, it's now almost orthodoxy to equate deontic constraints with agent-relative reasons:

¹⁸ An obvious objection to this point is that the agent-neutral counterpart reason generated by (9) *can*, in fact, be satisfactorily defined in terms of reason to do an act. Compare:

- (9) (x, φ) (It's being the case) that doing φ has a propensity to benefit x gives x reason to promote [i.e. *do*] φ .
- (12) $(x, \varphi) (\exists y)$ (It's being the case) that y 's doing φ has a propensity to benefit y gives x reason to do whatever promotes y 's doing φ .

Nevertheless, in so far as (12) directs a given agent to perform some *unspecified* action-type which promotes y 's doing φ (the specified action-type), the requirement of the reason cannot be satisfactorily defined in terms of giving an agent reason to directly perform the *action-type in question*. Again, if the reason were characterized in such a way, then the counterpart would not reflect the possibility of altruism. Recall:

- (6) $(x, \varphi) (\exists y)$ (It's being the case) that y 's doing φ has a propensity to benefit y gives x reason to φ .

I thank the anonymous reviewer for raising this concern.

¹⁹ It's worth reiterating that the dissociation argument was eventually abandoned (*Possibility*, p. vii; *View*, p. 159). I emphasize this because it reveals that the distinction was originally designed to undertake some substantive theoretical work, i.e. it formed the essential part of Nagel's argument for altruism. Within the latter context of *The View from Nowhere*, the idea that agent-relative reasons have to be subsumed under their agent-neutral counterparts is not explored. Indeed, the distinction is introduced simply as a distinction between the way reasons can vary in their *relativity to the agent*, where both agent-neutral and agent-relative reasons are characterized, *symmetrically*, in terms of reasons for anyone to *do* or *want* something (*View*, pp. 152–3).

Deontological constraints are agent-relative reasons which depend not on the aims or projects of the agent but on the claims of others. . . . If they exist, they restrict what we may do in the service of either relative or neutral goals.²⁰

However, despite this generally unquestioned orthodoxy, it remains unclear how these constraints are to be captured formally, for a distinct category of ‘negative reasons’ was never a part of Nagel’s original system:

We do not have to define a separate category of negative reasons, for any predicate *R* which provides a reason *against* the occurrence of anything to which it applies, will have corresponding to it another predicate *Q* which provides a reason *for* anything to which it applies, and which covers the same territory. We need only define *Q* as the predicate which holds of φ when *R* holds of the non-occurrence of φ .²¹

However, while this is true, it is unsatisfactory when trying to express the *force* of deontic constraints. If we want to understand deontic constraints qua agent-relative restrictions (i.e. agent-relative prohibitions on specific action-types open to agents), then we need to look further.

In his response to Skorupski, Broome suggested that Nagel would formulate constraints in the following way:

- (13) (x, φ) (It’s being the case) that φ is the telling of a lie by x gives x reason to promote the non-occurrence of φ .
- (14) (x, φ) (It’s being the case) that φ is the killing of an innocent person by x gives x reason to promote the non-occurrence of φ .
- (15) (x, φ) (It’s being the case) that φ is the breaking of a promise by x gives x reason to promote the non-occurrence of φ .

Following David McNaughton and Piers Rawling (and Portmore), the idea is that constraints are said to be *author agent-relative* in the sense that each agent has a special responsibility not to violate the constraint themselves – qua author of their own actions.²² Indeed, Nagel was keen to stress that ‘deontological reasons have their full force against *your* doing something – not just against it happening’.²³ These agent-relative constraints are then said to contrast with the

²⁰ Nagel, *View*, p. 175.

²¹ Nagel, *Possibility*, p. 45, n. 1.

²² David McNaughton and Piers Rawling, ‘Agent-Relativity and the Doing–Happening Distinction’, *Philosophical Studies* 63 (1992), pp. 167–85; ‘Value and Agent-Relative Reasons’, *Utilitas* 7 (1995), pp. 31–47; ‘Agent-Relativity and Terminological Inexactitudes’, *Utilitas* 7 (1995), pp. 319–25; ‘Deontology’, *The Oxford Handbook of Ethical Theory*, ed. David Copp (Oxford, 2009), ch. 15; Douglas W. Portmore, ‘Agent-Centred Restrictions’, *International Encyclopedia of Ethics*, ed. Hugh LaFollette (2013).

²³ Nagel, *View*, p. 187.

following agent-neutral counterparts, requiring agents to *minimize* the general occurrence of such constraint violation:

- (16) (x, φ) (It's being the case) that φ is the telling of a lie gives x reason to promote the non-occurrence of φ .
- (17) (x, φ) (It's being the case) that φ is the killing of an innocent person gives x reason to promote the non-occurrence of φ .
- (18) (x, φ) (It's being the case) that φ is the breaking of a promise gives x reason to promote the non-occurrence of φ .

Now, I'll argue below that agent-neutral reasons are not, inherently, utilitarian exhortations to minimize or maximize.²⁴ Nevertheless, the more serious problem here is that all of the agent-relative reasons above are equivalent to their agent-neutral counterparts because, for Nagel, *promoting* the non-occurrence of an event covers, primarily, the agent's own actions (recall, in the limiting case, if φ is an action, *doing* φ counts as *promoting* φ). And, as far as deontic constraints are concerned, their primary function is to prohibit the performance of certain *action-types* by the agent in question. Moreover, there is a clear difference between the principle 'It is wrong to tell lies' and the principle 'You should minimize the telling of lies'. Broome suggests that Nagel would capture the difference between these two principles *in terms* of the distinction between (author) agent-relative reasons and what we might call 'utilitarianized' agent-neutral reasons:

²⁴ Indeed, Nagel once declared: '[T]he principle of agent-neutrality does not automatically yield a species of utilitarianism, or some other counterintuitive principle as the method for deciding interpersonal conflicts. The requirement of agent-neutrality demands that full weight be accorded to the distinction between persons and to the irreducible significance of lives' (*Possibility*, p. 41). There may well be utilitarians who insist there is simply one agent-neutral reason for anyone and everyone to do what has a propensity to maximize aggregate utility. But the plausibility of such a position will need to be argued for independently – perhaps via the idea that practical rationality is instrumental and maximizing. Utilitarianism is not a corollary of the formal condition of agent-neutrality, nor is it an inherent feature of any kind of reason for action. Although Nagel is guilty of encouraging this reading by insisting in later works that agent-neutrality *is* inherently tied to the utilitarian's minimizing/maximizing rationale: 'Agent-neutral values are the values of certain occurrences or states of affairs which give everyone a reason to promote or prevent them. If murder is bad in an agent-neutral sense, for example, it means everyone has a reason to try and minimize the overall number of murders, independent of who commits them – and this might in some circumstances mean murdering a few people to prevent the murder of a large number. But if, on the other hand, murder is wrong in an *agent-relative* sense, this means each agent is required not to commit murder *himself*, and nothing is directly implied about what he must do to prevent murders by others' ('Personal Rights and Public Space', *Philosophy & Public Affairs*, 24 (1995), pp. 83–107, at 87–8). But, again, like agent-neutral reasons, there's nothing inherently utilitarian about agent-neutral *value*. Utilitarianism is found only in how one is disposed or required to respond to value. The idea that a value, itself, could be utilitarian is an odd one to say the least.

- (19) (x, φ) (It's being the case) that φ is the telling of a lie by x gives x reason to promote the non-occurrence of φ .
- (20) (x, φ) (It's being the case) that φ increases the number of lies told gives x reason to promote the non-occurrence of φ .

However, the reason expressed in (19) is not genuinely agent-relative; it's equivalent to:

- (21) (x, φ) (It's being the case) that φ is a lie gives x reason not to φ .

As Skorupski points out:

In general, any agent-neutral reason-predicate, ' $R\varphi$ ', can be converted into an agent-relative predicate, ' φ is an R -ing by x '. That does not show that it expresses an agent-relative reason. The question is whether you *can* express the reason with an agent-neutral predicate.²⁵

There is, then, a further tension between Nagel's formal dichotomy and the alleged agent-relativity of deontic constraints.²⁶ Indeed, Nagel openly admits that the agent-relative status of constraints is somewhat puzzling:

Deontological reasons have their full force against *your* doing something – not just against it happening. . . . You shouldn't break a promise or tell a lie for the sake of some benefit, even though you would not be required to forgo a comparable benefit in order to prevent someone else from breaking a promise or telling a lie. . . . The relative character cannot come simply from the character of the interest that is being respected, for that alone would justify only a neutral reason to protect the interest. And the relative reason does not come from an aim or project of the individual agent, for it is not conditional on what the agent wants. It is hard to understand how there could be such a thing. One would expect that reasons stemming from the interests of others would be neutral and not relative.²⁷

In an attempt to demystify this position, Nagel appeals to the doctrine of double effect: the idea that there is a significant distinction between *intentionally* bringing about an outcome and causing an outcome that is *foreseeable*, but not intentional:

²⁵ Skorupski, *Ethical Explorations* (Oxford, 1999), p. 53.

²⁶ McNaughton and Rawling (and Portmore) formulate an alternative version of the distinction specifically to accommodate the putative notion of *author* agent-relative constraints and utilitarianized agent-neutral reasons. In fact, the purpose of their reformulation is an attempt to demonstrate that Nagel (and Philip Pettit) cannot accommodate author agent-relativity. My own view is that author-relative reasons are not agent-relative reasons. Granted, one may wish to defend an author-relative rationale for the existence of constraints, but this will not align meaningfully with a non-utilitarianized (pre-theoretical) notion of an agent-neutral reason. There is already a distinction between deontological and utilitarian reasons, and it is *just that*.

²⁷ Nagel, *View*, p. 187.

The principle says that to violate deontological constraints one must maltreat someone intentionally. The maltreatment must be something that one does or chooses, either as an end or as a means, rather than something one's actions merely cause or fail to prevent but that one doesn't aim at.²⁸

The particular rationale for a deontological constraint, then, is that its violation amounts to the *intentional victimization* of the affected agent. However, this is hardly an appeal to (author) agent-relativity, in the sense that each agent is required not to commit intentional violations *himself*. Indeed, by moving away from the idea that constraints are creatures of the agent's perspective, Nagel moves away from the notion of an (author) agent-relative reason entirely:

[T]he victim feels outrage when he is deliberately harmed even for the greater good of others, not simply because of the quantity of harm but because of the assault on his value of having my actions guided by his evil. What I do is immediately directed against his good: it doesn't just in fact harm him.²⁹

Here, no appeal is made to the notion of an (author) agent-relative reason; there is simply the intentional assault on an individual's value which is a wrongdoing on the part of the agent.³⁰ In fact, Nagel ultimately abandoned the idea that the rationale for deontic constraints could be sought in the evil intentions of the victimizer, in favour of a rationale grounded solely in the inviolable status of those victimized:

The status is of a certain kind of *inviolability*, which we identify with the possession of rights, and the proposal is that we explain the agent-relative constraint against certain types of violation in terms of the universal but non-consequentialist value of inviolability itself.³¹

So, the underlying rationale for deontic constraints is no longer (author) agent-relative, yet, deontic constraints remain a species of agent-relative reason (in the author agent-relative sense).³² And, given

²⁸ Nagel, *View*, p. 179. Notice that the doctrine of double effect encompasses a different agency rationale than the (author) agent-relative account. The central feature of the author agent-relative account is that the requirement that one not violate a deontological constraint is more stringent than a demand to prevent violations. The significance of agency associated to the doctrine of double effect, however, is that it is impermissible for an agent to *intentionally* violate a constraint in order to secure some greater good. As it stands, the author agent-relative account of agency makes no appeal to intention (see Richard Brooke, 'Agency and Morality', *The Journal of Philosophy* 88 (1991), pp. 190–212).

²⁹ Nagel, *View*, p. 184.

³⁰ See Eric Mack, 'Deontological Restrictions Are Not Agent-Relative Restrictions', *Social Philosophy & Policy* 15 (1998), pp. 61–83.

³¹ Nagel, 'Personal Rights', p. 89; cf. David McNaughton and Piers Rawling, 'On Defending Deontology', *Ratio* 11 (1998), pp. 37–54.

³² This may be to avoid the so-called 'paradox of deontology': the idea that one is not permitted to violate a restriction even when its violation prevents further violations of

that author agent-relativity was never part of Nagel's original system, I suggest the theoretical notion of an author agent-relative reason – as something meaningfully distinguishable from a non-utilitarianized agent-neutral reason – should be abandoned.

VI. A TENSION IN SKORUPSKI'S ACCOUNT

Now that we're clear on the difficulties that beset the idea of author agent-relative constraints, we can isolate a different kind of tension within Skorupski's formulation of the distinction.

The purpose of Skorupski's original note was to challenge the tendency to make agent-relativity a necessary feature of deontological ethics. Nevertheless, while I certainly agree that agent-relativity is not the defining feature of deontological ethics – the defining feature is simply the claim that there are valid moral principles which do not depend on a prior theory of intrinsic value – there are some fundamental issues with Skorupski's formal dichotomy which need to be addressed.

To begin with, consider how Skorupski formulates a deontic constraint against the killing of innocent people:

- (22) (x, φ) (It's being the case) that φ is the killing of an innocent person gives x reason not to φ .

In this instance, the putative moral principle 'It's wrong to kill innocent people' entails an agent-neutral reason for action.³³ Likewise, the principle 'It's wrong to break promises' entails the following (in reasoning terms):

- (23) (x, φ) (It's being the case) that φ is the breaking of a promise gives x reason not to φ .

Now, this is where Skorupski's position begins to get confusing. In line with the reasons not to kill innocent people and not to break promises, Skorupski once proclaimed the reason there is not to tell lies is also agent-neutral:

- (24) (x, φ) (It's being the case) that φ is a lie gives x reason not to φ ,³⁴

the same type by others (or oneself). But this idea is paradoxical only if the advocate of deontic constraints accepts that the disvalue or objectionable property of their violation is to be minimized in the first place.

³³ Skorupski, *Explorations*, p. 52.

³⁴ Skorupski, *Explorations*, p. 52.

affirming elsewhere that the principle ‘It’s wrong to tell lies’ is agent-neutral in the following sense:

The agent-neutral reason-predicate is ‘It is wrong to tell lies and φ is a lie’. In other words, I am taking it that (x) (It is wrong to tell lies and φ is a lie \rightarrow x has reason not to do φ).³⁵

However, in a footnote that occurred later in *Ethical Explorations*, Skorupski U-turns on the status of the reason not to tell lies without reference to his earlier view, declaring:

If there’s reason not to tell a lie, just because it’s a lie, and ‘lying’ means ‘saying something one believes to be false’, then that reason is *agent-relative*.³⁶

(25) (x, φ) (It’s being the case) that φ is saying something x believes to be false gives x reason to φ .

Furthermore, in his most recent discussion of the distinction, the status of the reason one has to *keep* one’s promises is explicitly identified as agent-relative. For instance, Skorupski states that the fact that he has promised to send Mira a book is an agent-relative reason for him to send her one, expressible via the following universal predicate:

(26) (x, φ) (It’s being the case) that x has promised to φ gives x reason to φ ,

adding:

This reason *cannot* be expressed by an agent-neutral predicate: that is one which contains no free occurrence of the relevant agent-denoting term.³⁷

There is, then, an ambiguity within Skorupski’s formulation of the distinction. In earlier works, reasons not to kill innocent people, break promises, tell lies, etc., are explicitly identified as agent-neutral. Yet, more recently, reasons not to tell lies, and reasons to keep one’s promises have been formally identified as agent-relative.

VII. QUALIFYING SKORUPSKI’S ACCOUNT

To understand why these ambiguities occur, we need to get clear on the differences between Skorupski’s predicate-based distinction, and

³⁵ Skorupski, ‘Neutral Versus’, p. 241, n. 6.

³⁶ Skorupski, *Explorations*, p. 60, n. 23.

³⁷ Skorupski, *Domain*, p. 64.

his most recent statement of the distinction formulated in terms of *essentially indexical* and *non-indexical* reason-giving facts.³⁸

[A] reason-giving fact is *essentially* indexical when the actor, in order to grasp it as a reason for *him*, or a reason for him to φ at some *time* identified by its relation to the present, must grasp it indexically as a fact about himself or about the present; that is when its indexicality is essential to its reason-giving force. . . . A reason is relative if and only if it is essentially indexical. . . . When the reason giving force of the fact, for the actor, does not rest on its indexicality we shall say the reason is *neutral*.³⁹

Defined in this way, agent-relative reasons are said to be indexical facts, where an indexical fact is simply ‘a true proposition expressible by a declarative sentence containing indexicals’.⁴⁰ From the first-person perspective, then, the fact that eating healthily has a propensity to benefit *me* is an agent-relative reason for *me* to eat healthily; the fact that *I* have promised to send Mira a book is an agent-relative reason for *me* to send her one; the fact that taking *my* grandfather fishing promotes his interests is an agent-relative reason for *me* to take him fishing. In these cases, an indexical must appear in my representation of the reason in order for me to recognize it as a reason for me to act at a given time.⁴¹ On the other hand, the fact that any contribution to some particular charity at any time will be used effectively by that charity for famine relief is considered to be an agent-neutral reason for me to contribute to charity. This fact need not be presented to me in an indexical way in order for me to recognize the reason it is for *anyone* to contribute to charity.

Skorupski does not address the subject of deontic constraints in relation to this fact-based version of the distinction. Nevertheless,

³⁸ Only recently has Skorupski identified the fact-based distinction and the predicate-based distinction as ‘equivalent’ alternatives (*Domain*, pp. 61, 63). In earlier works, the predicate-based distinction and the fact-based distinction are not explicitly distinguished (*Neutral Versus*, pp. 244–8; *Explorations*, ch. 3).

³⁹ Skorupski, *Domain*, pp. 60–1. The ontological status of these facts needn’t concern us, though we can note that Skorupski subscribes to a ‘nominal’ conception of reason-giving facts, taking them to be *sui generis*, abstract entities standing in a ‘one–one’ relation to true propositions, rather than ontologically substantive entities qua obtaining states of affairs (*Domain*, pp. 61–3; cf. Rønnow-Rasmussen, *Personal*, ch. 9; ‘Reasons and Two Kinds of Fact’, *Discusiones Filosóficas* 13 (2012), pp. 95–113; cf. Buckland, ‘Normative Reasons’).

⁴⁰ Skorupski, *Domain*, p. 62. Philip Pettit draws the distinction in a similar way. The idea is that agent-relative reasons are propositions that essentially involve *pronominal back-reference* to the agent for whom they are reasons. In each instance, the motivating consideration (the reason-giving fact or principle) involves essential reference to *him* or *his* (‘Universalizability without Utilitarianism’, *Mind* 96 (1987), pp. 74–82, at 75; ‘The Paradox of Loyalty’, *American Philosophical Quarterly* 25 (1988), pp. 163–171, at 165).

⁴¹ The same is applicable in third-person cases.

regarding the reason not to tell lies, the earlier idea was that the indexicality of the reason is essential to its rationalizing force:

If there's reason not to tell a lie just because it's a lie, and 'lying' means 'saying something one believes to be false', then that reason is agent-relative. One doesn't have to know 'who one is' to grasp that one has it in a particular case, but one still needs to know *something* about oneself, viz., that one believes some particular thing to be false. This indexical fact is open to reflective awareness, but the point remains that there is an indexical fact about oneself that one needs to know.⁴²

The case is said to be different when dealing with the reason to give money to famine relief, or the reason not to kill innocent people:

I don't have to know any such indexical fact, about who I am, what relations I stand in, or what time it is now, to know that if it's obligatory to contribute to famine relief, there's reason – reason for me now – to contribute to famine relief; or that if it's wrong to kill the innocent, there's reason for me now not to kill the innocent. . . . I don't have to know who I am, or any other fact about me to know that it's wrong to kill innocent people, and hence that I have reason not to kill innocent people. That's because this is *not* an agent-relative reason.⁴³

Now, I don't want to deny that saying something one believes to be false gives one reason not to say it – because it's lying and, just like killing innocent people, it is, *prima facie*, wrong to tell lies. Nevertheless, this reason is not agent-relative in the same sense as, say, the reason there is not to neglect one's talents, or the reason there is not to neglect one's family and friends, or the reason there is to keep one's promises. These reasons are *object agent-relative* in the sense that they're relativized to specific individuals given their histories and a significant relationship to an object of moral concern; much like your child, your promise is an entity that you have previously 'created'. A reason not to tell lies is not object agent-relative in this sense. Granted, Skorupski admits that one doesn't need to know 'who one is' in order to know that there's reason for them not to tell lies. However, it's *not* the case that an agent must believe some particular thing to be false in order to grasp that there's reason for them not to tell lies at a given time. 'Lying' may well mean 'saying something one believes to be false', but it's perfectly conceivable that I could know that it's wrong to tell lies, and that there's reason for me, now, not to tell lies – *because I know it's wrong to tell lies* – without having to believe some particular thing to be false. I know, for instance, that it would be wrong for me to lie to my girlfriend about my whereabouts on Friday evening, but there's no particular thing I need to believe is false in order for me to know that it's wrong for me to lie

⁴² Skorupski, *Explorations*, p. 60, n. 23.

⁴³ Skorupski, *Explorations*, p. 60.

to her, qua acknowledgement of the putative agent-neutral principle 'It's wrong to tell lies', or even acknowledgement of the putative agent-relative principle 'It's wrong to lie to one's girlfriend'.

Far from qualifying things, this leaves us in a bit of a mess. Granted, Skorupski's focus remains the same, i.e. emphasis on determining the agent-neutral/relative status of a reason depends on whether or not the reason-giving predicate *can* be expressed agent-neutrally. But, as we've just seen, the machinery can generate conflicting results depending on whether we appeal to reason-giving facts or reason-giving principles.

VIII. CAPTURING AGENT-RELATIVITY

How, then, should we capture genuine agent-relativity? My proposal is that there is an inherent connection between *object agent-relativity* and the essentially indexical account of agent-relative reasons, which does not hold in the case of *author agent-relativity*, where emphasis is placed on whether the reason can be meaningfully transposed without indexical reference to the agent without redundancy. There is, then, a significant difference between, say:

- (27) (x, y) That y is x 's child is a reason for x to care for them (that is y).
- (28) (x, y) That y is x 's promise is a reason for x to keep it (that is y).
- (29) (x, y) That y is x 's ambition is a reason for x to pursue it (that is y).

And:

- (30) (x, φ) That φ is x 's lie is a reason for x not to tell it (that is φ).
- (31) (x, φ) That φ is x 's killing of an innocent person is a reason for x not to do it (that is φ).
- (32) (x, φ) That φ is x 's contribution to charity is a reason for x to do it (that is φ).

On this view, (27), (28) and (29) are *genuinely* agent-relative because they cannot be meaningfully transposed without indexical reference to the agent without redundancy. On the other hand, the author agent-relative (30), (31) and (32) *can* be meaningfully transposed without indexical reference to the agent without redundancy:

- (33) (x, φ) That φ is a lie is a reason for x not to tell it (that is φ)
- (34) (x, φ) That φ is the killing of an innocent person is a reason for x not to do it (that is φ)
- (35) (x, φ) That φ is a contribution to charity is a reason for x to do it (that is φ)

Those reasons which are object agent-relative are genuinely agent-relative. Those reasons which are not object agent-relative are agent-neutral.

IX. AGENT-RELATIVITY AND THE STATUS OF DEONTOLOGY

To be clear, my claim is not that author relativity is incoherent *as an account of the underlying rationale for deontic constraints* – certainly, the author rationale for constraints is one that needs careful consideration. My claim is, rather, that the theoretically substantive notion of *author agent-relativity* cannot be distinguished from the pre-theoretical notion of *non-utilitarianized* agent-neutrality in any meaningful way, and does not, therefore, constitute a genuine distinction between agent-neutral and agent-relative reasons for action.⁴⁴ The genuine distinction is simply a distinction between those reason-giving considerations (be they facts or putative moral principles) which encapsulate an essentially indexical relationship between an agent and an object of moral concern, and those which do not.

This observation lends support to Skorupski's idea that the tendency to make agent-relativity a necessary feature of deontological ethics is mistaken, i.e. that no thesis about agent-neutrality versus agent-relativity should be a defining feature of deontological ethics.⁴⁵ To see why, consider that 'teleology' is understood, fundamentally, as a theory of desired *ends* – what's good. Pure teleological ethics, then, affirms that deontic verdicts are derived from a prior theory of intrinsic value. In this sense, there can be both agent-neutral and agent-relative teleological ethical theories. Utilitarians share the same, fundamental agent-neutral principle of action: 'Everyone has reason to do what has a propensity to maximize aggregate utility'. This is why utilitarianism can be categorized as an agent-neutral, teleological ethical theory; it gives all agents the same substantive and ultimate aim of maximizing aggregate utility. Egoists, on the other hand, share

⁴⁴ I've not discussed McNaughton and Rawling's version of the distinction in great detail here (nor Portmore's). However, I will say that any reason one might have for favouring their version of the distinction over the distinction I have outlined above will be circular: in order to give preference to their version of the distinction you must *first* accept that the rationale for deontic constraints is both author agent-relative and agent-focused. In this sense, their distinction amounts to a theoretically loaded distinction between author agent-relativity and utilitarian agent-neutrality. And, as we've already seen, author agent-relativity is a controversial notion in itself, and there is nothing *inherently* utilitarian about agent-neutrality.

⁴⁵ Skorupski, 'Neutral Versus', p. 242; *Explorations*, p. 56; cf. McNaughton and Rawling, 'Inexactitudes'; Portmore, *Commonsense Consequentialism* (Oxford, 2013).

the same fundamental agent-relative principle of action: ‘Everyone has reason to do what has a propensity to benefit themselves’. This gives each agent the same substantive and ultimate aim of maximizing their own utility. In this sense, egoism can be categorized as an agent-relative, teleological ethical theory.⁴⁶

What, then, of ‘deontology’? Following Skorupski, deontology is best understood as a theory of *duty* – a theory that concerns itself with (morally) right or wrong action. *Pure* deontological ethics, then, is an ethics that does not depend on a prior theory of intrinsic ethical value, i.e. what’s good or bad (morally speaking). Skorupski takes Kant’s ethics to be *purely* deontological in this sense:

[T]he concept of good and evil is not defined prior to the moral law, to which, it would seem, the former would have to serve as foundation, rather the concept of good and evil must be defined after and by means of the law.⁴⁷

Indeed, in light of this idea we can see that debates concerning whether Kantian moral prohibitions must be understood as agent-relative – even if we allow that all reasons are teleological – are misguided.⁴⁸ The categorical imperative itself (qua fundamental principle of morality) is deontological and agent-relative:

Act only in accordance with that maxim through which *you* can at the same time will that it become a universal law.⁴⁹

⁴⁶ Notice that this goes against the traditional view that agent-neutral theories give agents *shared* aims, whereas agent-relative theories give agents *distinct* aims. In both cases, our *reasons* may vary; after all, what maximizes my utility may not maximize yours. Nevertheless, qua egoists we both share the same substantive and ultimate agent-relative aim of maximizing our own utility.

⁴⁷ Immanuel Kant, *Critique of Practical Reason*, trans. Lewis White Beck (New York, 1956), p. 65; Skorupski, *Explorations*, p. 55.

⁴⁸ Cf. Ridge, ‘Neutral vs. Relative’; cf. Vaughn Huckfelt, ‘Categorical and Agent-Neutral Reasons in Kantian Justifications of Morality’, *Philosophia* 35 (2007), pp. 23–41.

⁴⁹ Skorupski has suggested that the categorical imperative is an agent-neutral principle because the end it posits is rational nature (‘Autonomy and Impartiality: *Groundwork* III’, *Kant’s Groundwork of the Metaphysics of Morals: A Critical Guide*, ed. Jens Timmerman (Cambridge, 2009), ch. 9, p. 176). Now, we might agree that each agent’s end is rational nature. However, as far as Kant is concerned, it’s necessary to stress that the agent is specifically concerned with acting in accordance with those maxims which she can will to become universal laws. In this instance, such reference to the agent’s willing is far from trivial – it is the object of moral concern. Kant’s concern is with the autonomy of the practical agent, and their ability to act on those reasons which they, qua particular individual, have first-personally endorsed. However, we should also note the fact that a maxim can be willed as a universal law just means it is permissible, and, intuitively, the fact that something is permissible entails only the absence of a reason – that there is no decisive moral reason to refrain from performing it. Moral permissibility alone entails no positive reason to perform the action.

Nevertheless, the subsidiary maxims through which one can act can take either agent-neutral or agent-relative form. The point is that many moral maxims are agent-relative, and give rise to agent-relative reasons: 'It's wrong to neglect *one's* children', 'It's wrong to break *one's* promises', etc. But there are also agent-neutral duties arising from principles such as 'It's wrong to kill innocent people', 'It's wrong to tell lies', 'Everyone should help those in need', etc. Ultimately, then, the deontological distinction between right and wrong action has nothing to do with the distinction between agent-neutral and agent-relative reasons. Consequently, it's a mistake to characterize deontological ethics solely in terms of agent-relativity.

X. CONCLUSION

I began this article by mediating the exchange that occurred between Skorupski and Broome concerning how best to formulate Nagel's distinction between agent-neutral and agent-relative reasons for action. Mediating this exchange revealed that Nagel's distinction contained an asymmetry in the way it was formulated, and that we do best by speaking in terms of reasons for agents to perform action-types open to them, rather than reasons for agents to promote the occurrence of events.

After clearing up some ambiguities with the Nagelian distinction and the status of deontic constraints, it emerged that there are, in fact, two notions of agent-relativity in play: *author agent-relativity* and *object agent-relativity*. I then argued that Nagel – the original proponent of agent-relative constraints – does not provide an adequate agent-relative rationale for their existence; Nagel's distinction is drawn solely in terms of object agent-relativity, and the notion of author agent-relativity is imported into his later account of deontic constraints without justification. From here, I maintained that object agent-relativity is the genuine notion of agent-relativity which, unlike author agent-relativity, can be mapped onto the essentially indexical account of agent-relative reasons without redundancy. Consequently, the (author) agent-relativity traditionally associated with deontic constraints is not genuine agent-relativity, and cannot be distinguished from a non-utilitarianized conception of agent-neutrality in any meaningful way. Again, the genuine distinction between agent-neutral and agent-relative reasons is simply a distinction between those reason-giving considerations (be they facts or putative moral principles) which encapsulate an essentially indexical relationship between an agent and an object of moral concern, and those which do not. This captures agent-relativity solely in terms of object agent-relativity, and lends definitive support to Skorupski's view that no

thesis about agent-neutrality versus agent-relativity should be a defining feature of deontological ethics.⁵⁰

jamie.buckland@york.ac.uk

⁵⁰ I am grateful to Richard Rowland and the two anonymous referees for *Utilitas* for extremely helpful comments on earlier versions of this article. Additionally, thanks are owed to Jonathan Dancy, Brad Hooker, Max de Gaynesford, Philip Stratton-Lake and Alan Thomas for helpful discussion on work from which this article was derived.