Agent-Relativity and the Status of Deontological Restrictions

Jamie Buckland¹

Accepted: 27 March 2021
© The Author(s) 2021

There is a familiar distinction in moral philosophy between deontological and consequentialist normative theories. Deontological theories place a significant emphasis on the idea that what matters, morally, is that individual agents perform or refrain from performing certain action-types, rather than promoting the occurrence, or non-occurrence of these action-types more generally. More specifically, a deontologist will posit the existence of various *deontological constraints* (or *side-constraints*) that prohibit (*deontological restrictions*), require (*deontological obligations*), or permit (*deontological permissions*) the performance of particular action-types because they are inherently right or wrong, over and above a consequentialist requirement to minimise or maximise the occurrence of those action-types because of ‘the Good’ the occurrence or non-occurrence of those action-types brings about (presuming such requirements even exist).

The distinction between deontological and consequentialist normative theories is straightforward enough. Nevertheless, *pace* Thomas Nagel and David McNaughton

✉ Jamie Buckland
jamie.buckland@york.ac.uk

¹ Department of Philosophy, University of York, York YO10 5DD, UK

and Piers Rawling, there is a long-standing project which seeks to demarcate the boundaries of deontology and consequentialism even further by defining deontology exclusively in terms of its *agent-relativity*.¹ Building on Derek Parfit's idea that agent-relative theories give each agent different aims (whereas agent-neutral theories give each agent a common aim), McNaughton and Rawling sought to capture a formal distinction between agent-neutral and agent-relative rules in terms of the variable aims they give to agents.² These rules can then be mapped to their corresponding normative theories, rendering all deontological normative theories agent-relative theories.³ Indeed, Matthew Hammerton, the most recent defender of what I will call McNaughton and Rawling's 'demarcation project' goes as far as stressing that deontology is *necessarily* agent-relative.⁴

The aim of this paper is to show that the demarcation project is problematic, and to defend both the logical possibility as well as the plausibility of an agent-neutral account of deontological restrictions. The success of the demarcation project (and the agent-relativisation of deontology) relies on drawing a formal distinction between agent-neutral and agent-relative rules which seeks to accommodate and contrast two theoretically substantive (and contentious) intuitions. First, a predetermined *agent-focused* (or *agent-centred*, or *agent-relative*) account of deontology where the numerical identity of the agent makes a difference to the deontic status of the act to the extent that agents have a particular interest in what *they do* as opposed to what they merely allow to *happen* (in contrast to a *victim-focused* (or *victim-centred*, or *victim-relative*) account of deontology grounded in the inviolable status of the would-be victim). And second, a predetermined agent-neutral account

¹ Thomas Nagel, *The View from Nowhere* (Oxford: Oxford University Press, 1986), Ch. XI; David McNaughton and Piers Rawling, "Agent-Relativity and the Doing-Happening Distinction", *Philosophical Studies*, Vol. 63, No. 2 (1991), pp. 167-185; David McNaughton and Piers Rawling, "Honoring and Promoting Values", *Ethics*, Vol. 102, No. 4 (1992), pp. 835-843; David McNaughton and Piers Rawling, "Deontology and Agency", *The Monist*, Vol. 76, No. 1 (1993), pp. 81-100; David McNaughton and Piers Rawling, "Value and Agent-Relative Reasons", *Utilitas*, Vol. 7, No. 1 (1995a), pp. 31-47; David McNaughton and Piers Rawling, "Agent-Relativity and Terminological Inexactitudes", *Utilitas*, No. 7, Vol. 2 (1995b), pp. 319-325; David McNaughton and Piers Rawling, "On Defending Deontology", *Ratio*, Vol. 11, No. 1 (1998), pp. 37-54; David McNaughton and Piers Rawling, "Achievement, Welfare, and Consequentialism", *Analysis*, Vol. 61, No. 2 (2001), pp. 156-162; Matthew Hammerton, "Distinguishing Agent-Relativity from Agent-Neutrality", *Australasian Journal of Philosophy*, Vol. 97, No. 2 (2019), pp. 239-250.

² Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), p. 27.

³ The converse does not hold, i.e., not all consequentialist theories are agent-neutral (cf. Desheng Zong, "Agent-Neutrality is the Exclusive Feature of Consequentialism", *Southern Journal of Philosophy*, Vol. XXXVIII (2000), pp. 677-693). McNaughton and Rawling deployed the distinction between agent-neutrality and agent-relativity in order to demarcate the line between what they referred to as 'deontological' and 'consequentialist' normative theories. However, their use of the term 'consequentialism' centrally describes classical, direct act-consequentialism whereby one's moral duty is to maximise *the good* (generally conceived). Consequentialism is now understood in a boarder sense to include forms of agent-relative consequentialism whereby one's moral duty is to maximise *one's own good*.

⁴ Matthew Hammerton, "Is Agent-Neutral Deontology Possible?", *Journal of Ethics and Social Philosophy*, Vol. 21, No. 3 (2017), pp. 319-324.

of consequentialism tied to an instrumental account of practical rationality where agent-neutral rules are defined as injunctions to either minimise or maximise.⁵

Alternatively, the pre-theoretical distinction between agent-neutral and agent-relative rules I defend in section 8 is free from these substantive considerations and is remarkably easy to articulate: namely, a distinction between universal rules which have either a restricted or a wide scope of practical application.

1 Agent-Relative Deontological Restrictions

Following McNaughton and Rawling, Hammerton draws a formal distinction between moral rules corresponding to the variable aims they give to individual agents:

A rule expressing a moral requirement is agent-relative *iff*, when the rule is represented with the form ‘(x) (x must {verb} that [...])’, there is an ineliminable occurrence of ‘x’ in the square brackets that is bound by the initial universal quantifier.

A rule expressing a moral requirement is agent-neutral *iff* it is not agent-relative.⁶

Take, for instance, a deontological restriction that prohibits anyone and everyone from killing innocent people (even if doing so is the only way to prevent more killing of innocents by others), alongside its agent-neutral (consequentialist) counterpart requiring agents to minimise the killing of innocent people:

(1) *Each agent must not kill innocent people:*

(x) (x must act such that [x does not kill innocent people])

(2) *Each agent must minimise the killing of innocent people:*

⁵ It is important not to confuse a *victim-focused* rationale for deontological restrictions with what Hammerton has referred to as ‘*patient-relativity*’ in morality (Matthew Hammerton, “Patient-Relativity in Morality”, *Ethics*, Vol. 27 (2016), pp. 6-27). As will become clear, appealing to a victim-focused rationale to underwrite deontological restrictions *grounds* the restriction in the inviolable status of those victimised (*within* the victim, as it were, as opposed to *within* the would-be victimiser). Patient-relativity, on the other hand, is when a moral patient’s numerical identity is said to make a difference to the deontic status of the act. For example, Hammerton takes the humanity formula of Kant’s categorical imperative – *Always treat people as ends in themselves and never as mere means* – as a possible theoretical basis for patient-relativity to the extent that it prohibits breaking a promise to a promisee now in order to keep several promises to other promisees in the future because it treats the present promisee (patient) as a ‘mere means’ to fulfilling promises owed to others. Nevertheless, while Kant’s humanity formula of the categorical imperative offers an interesting and welcome theoretical basis for patient-relativity in morality, to the extent that it holds that each agent must act such that *they* treat people as ends in themselves and never as mere means, it is, at base level, on Hammerton’s account, still an agent-relative rule giving different agents the different ultimate aim of not breaking a promise to a particular promisee *now*.

⁶ Hammerton, “Distinguishing Agent-Relativity from Agent-Neutrality”, p. 247.

(x) (x must bring it about, to the greatest degree that x can, that [there is minimal killing of innocent people])

This formalisation of the distinction is said to be attractive because it allows us to make sense of Parfit's idea that moral rules in ordinary language can be classified as either agent-neutral or agent-relative in virtue of the differing aims they give to agents:

[A]ll rules can be translated such that agents must perform a type of act (such as ensuring, acting such that, bringing it about) with regard to some propositional content. [...] [W]hen the relevant variable occurs in a rule's content, this corresponds to that content being different for different agents, whereas the absence of the relevant variable corresponds to the content being the same for all agents. [...] If each rule's content is regarded metaphorically as what the rule is *aimed at*, then our formal account makes sense of the idea that some rules (but not others) give different aims to different agents.⁷

The central idea behind this distinction is the thought that 'some moral rules (or reasons, or theories) give agents a special focus on themselves, their actions, or states of affairs concerning them, whereas other rules lack this personal focus, and instead give agents a general concern with actions and states of affairs that involve anyone'.⁸ Direct-rule (1) is said to be an agent-relative deontological rule in virtue of the fact that it gives each agent a different, *ultimate aim*. It gives *you* the aim that *you* don't kill innocent people, and it gives *me* the aim that *I* don't kill innocent people. Minimising-rule (2), on the other hand, is said to be an agent-neutral consequentialist rule in virtue of the fact that it gives each agent the same, ultimate aim. It gives *you* the aim that *you* ensure that there is minimal killing of innocent people in the world, and it gives *me* the aim that *I* ensure that there is minimal killing of innocent people in the world.

These results are then said to generalise; the broader class of deontological constraints (restrictions, special obligations, and permissions) are all, necessarily, agent-relative constraints because they give each agent a different ultimate aim, and each deontological constraint is said to have an agent-neutral, consequentialist analogue. Hammerton takes this formal distinction between agent-neutral and agent-relative rules to be 'the best formal account of the distinction' available.⁹

Before moving on, it is worth noting that there is nothing '*metaphorical*' in our understanding of what rules (1) and (2) are *aimed at*. Quite clearly, direct-rule (1) states that *each agent is required to act such that they do not kill innocent people*, i.e., *each agent must bring about a world in which they don't kill innocent people*, whereas minimising-rule (2) states that *each agent must bring it about, to the greatest degree that they can, that there is minimal killing of innocent people*, i.e., *each agent must bring about a world in which there is minimal killing of innocent*

⁷ Hammerton, "Distinguishing Agent-Relativity from Agent-Neutrality", p. 248.

⁸ Hammerton, "Distinguishing Agent-Relativity from Agent-Neutrality", p. 239-240.

⁹ Hammerton, "Distinguishing Agent-Relativity from Agent-Neutrality", p. 239.

people, *simpliciter*. Moreover, it is somewhat artificial to maintain that each agent must *act such* that they don't kill innocent people. There is a good chance I could, albeit accidentally, kill an innocent person whilst cycling to work tomorrow, but I am not required to avoid cycling for this reason, and I am not justified in telling my boss that I won't be at work tomorrow (or for the foreseeable future) for this reason. The intuitive thought concerning deontological restrictions explicitly concerns *inaction*, i.e., the thought that agents are prohibited from performing token action-types (because the performance of those action-types is inherently wrong).¹⁰ In light of this, we do better to speak of a prohibition on murder *qua* the intentional killing of a fellow agent (be they innocent or not) as a paradigm example of a deontological restriction.¹¹

2 Agent-Neutral Deontology?

Hammerton's view is that deontological restrictions are, *necessarily*, agent-relative restrictions. Nevertheless, Tom Dougherty has defended the possibility of agent-neutral deontology by positing the existence of agent-neutral deontological restrictions within Hammerton's logical space.¹²

Consider the following rule:

(3) *Each agent should ensure that no one murders to prevent more murdering by others:*

(x) (x should ensure that [no one murders to prevent more murdering by others])

Rule (3) is *both* agent-neutral *and* deontological. It is agent-neutral in the sense that it gives all agents the same, ultimate aim of ensuring that *no one* murders to prevent more murdering by others, and it is deontological in the sense that it posits a universal restriction on preventive murdering; indeed, in certain situations it requires that the agent does not murder (even when they have an opportunity to murder to prevent more murdering by others). It does not, however, give the agent any special concern with their own murdering *qua authors of their own action*. Certainly, it requires the agent to ensure something is true of them, but only insofar as they are one agent among many.

This is a welcome addition to the realm of deontological theory; it captures the idea that all moral agents can share a unified 'moral vision' – a deontology grounded

¹⁰ I am thinking of Kant's idea that the concept of good and evil is not defined prior to the moral law, i.e., a 'pure' deontology that does not depend on a prior theory of intrinsic value (Immanuel Kant, *Critique of Practical Reason*, trans. Lewis White Beck (New York, 1965), p. 65; John Skorupski, *Ethical Explorations* (Oxford: Clarendon Press, 1999), p. 55; Jamie Buckland "Skorupski and Broome on the Agent-Neutral/Relative Distinction", *Utilitas*, Vol. 31, No. 1 (2019), pp. 59-82.

¹¹ I'll return to these points about aims and metaphor below.

¹² Tom Dougherty, "Agent-Neutral Deontology", *Philosophical Studies*, Vol. 162, No. 2 (2013), pp. 527-537. Dougherty's paper is directed towards McNaughton and Rawling's version of the distinction. I have updated (3) to accord with Hammerton's 2019 version.

in the idea that everyone share the same goals and preferences regarding the performance of particular action-types, rather than a deontology grounded in an agent's desire to avoid 'dirty hands'.¹³

3 Hammerton's Response: Deontology is Necessarily Agent-Relative

Notwithstanding the attraction of such a position, Hammerton has argued that agent-neutral deontology is impossible; deontological restrictions cannot be *both* agent-neutral (in the sense that they give all agents the same ultimate aim) *and* deontological (in the sense that they remain restrictions). So, while (3) can be understood as *either* agent-neutral or deontological, there are no interpretations of (3) in which it is *both* agent-neutral and deontological.¹⁴

In expounding this idea, the first thing to note is that, on a literal interpretation, (3) requires each agent to guarantee that no one murders to prevent more murdering by others, which is extremely demanding. This would require an agent to ensure that there are no preventive murders even when they lacked the opportunity or the ability to do so. Moreover, insofar as a preventive murder were to occur, then everyone (including those unable to prevent the murder) would be in the wrong for not having prevented it.

In light of these difficulties, Hammerton reinterprets (3) in a manner that does not place unrealistic and excessive demands on an agent:

(4) *Each agent, whenever they have the ability and the opportunity to do so, should ensure such that no one murders to prevent more murdering by others:*

(x) (whenever x has the ability and the opportunity, x must ensure that [no one murders to prevent more murdering by others])

Nevertheless, (4) gives rise to further issues.

Consider the following case:

Preventive Murder: Tom and Mary are each about to commit a preventive murder. The only way for you to stop them is by committing one preventive murder yourself (perhaps the horror of witnessing a preventive murder will cause them to abandon their plans).

In this scenario, you don't have the ability to ensure that no one commits a preventive murder, so you are not in the wrong should you fail to prevent Tom and Mary from preventive murdering. Nevertheless, you are presented with a moral quandary: either commit one preventive murder yourself, or allow two murders to occur. But, if (4) were a genuine deontological restriction, it would prohibit *you* from preventive murdering, even if doing so were the only way to prevent more preventive murders

¹³ Dougherty "Agent-Neutral Deontology", pp. 531-533.

¹⁴ Hammerton, "Is Agent-Neutral Deontology Possible?"

by others, as in (1), i.e., there would be no quandary. But there *is* a quandary, and (4) gives you no guidance on how to resolve this quandary. So, while the modified (4) is an agent-neutral rule, it is *not* a deontological restriction.

We could attempt to get around this issue by offering a further modification:

(5) *Each agent should, to the best of their ability, ensure that no one performs a preventive murder, making sure that whenever they cannot prevent all such murders, they at least do not perform any themselves:*

(x) (x should, to the best of x's ability, ensure that [no one performs a preventive murder, making sure that whenever x cannot prevent all such murders, x at least does not perform any])

However, while (5) is a deontological restriction, it is not an agent-neutral rule. Likewise, with (6):

(6) *Each agent should, to the best of their ability, ensure that no one performs a preventive murder, minimising the total number of such murders whenever they cannot prevent them all:*

(x) (x should, to the best of x's ability, ensure that [no one performs a preventive murder, minimising the total number of such murders whenever x cannot prevent them all])

While (6) gives all agents the same, ultimate aim that there are as few preventive killings as possible, it does not, as a deontological restriction must, prohibit you from murdering.

Ultimately, then, no reading of (4) can be agent-neutral while, simultaneously, giving a deontological verdict in cases like *Preventive Murder*. Hence, Hammerton's insistence that deontology is, necessarily, agent-relative.

4 A Circular Necessity

At this point we might be tempted to think that Hammerton has dispensed with the notion of agent-neutral deontology. Nevertheless, the rationale for supporting this conclusion is both circular and contentious: in order to accept Hammerton's conclusion, you must *first* accept that deontology is necessarily an agent-centred (or agent-relative) enterprise, i.e., you must first accept that the underlying rationale for deontological restrictions is agent-relative in the sense that deontology is *grounded in* the idea that agents have a special responsibility to refrain from performing certain action-types themselves *qua authors of their own actions*.¹⁵ But these are

¹⁵ It is important to be clear what I mean by placing emphasis on the '*qua authors of their own action*' locution. No one will deny that, as far as restrictions are concerned, everyone has a special responsibility not to violate them; the act that constitutes the violation will be *my* act. Nevertheless, the position I challenge below is that the restriction is *grounded in* and *agent-relativized via* the idea that agents are subject to restrictions *qua authors of their own actions*.

theoretically substantive ideas, and should not be built into the notion of deontology from the outset.

To highlight the circularity in question, consider how an agent-neutral consequentialist rule directing agents to minimise murdering in the world can take both agent-neutral and agent-relative forms via Hammerton's machinery without a change in content:

(7) (x) (x must bring it about, to the greatest degree that x can, that [there are minimal murders in the world])

(8) (x) (x must act such that [x minimises murders in the world])

Each of these of rules is analytically equivalent – *a priori*, there is no possible world in which one is fulfilled, and the other is violated. In light of the possibility of this kind of transformation, Hammerton needs some way of identifying those rules which are *genuinely* agent-relative, and those which are imposters (generated by a triviality). He achieves this via a direct appeal to the moral significance of the agency relation, i.e., the idea that, where deontological restrictions are concerned, agents have a special responsibility to refrain from performing action-types themselves *qua authors of their own actions*. Indeed, as far as the direct-rule prohibiting each agent from murdering is concerned, the rationale underwriting its agent-relativity is that each agent has a special responsibility to *act such* that they don't murder people themselves – *qua authors of their own action* – as opposed to ensuring there is minimal murdering *occurring* – even if that means carrying out a preventive murder oneself.¹⁶ It is this notion of '*author agent-relativity*' (as McNaughton and Rawling originally referred to it) that is said to underwrite deontological restrictions, as opposed to accounts which appeal to the status of the victim, or the relationship between the agent and victim. Consequently, Hammerton's reformulation of the distinction is not simply a syntactic distinction designed to demarcate the boundaries of deontology and traditional, direct act-consequentialism (the former being necessarily agent-relative and the latter being agent-neutral). Rather, the distinction is constructed to accommodate substantive assumptions about the nature of deontology, i.e., a specifically agent-centred account of deontology underwritten by an account of author agent-relativity (as opposed to a victim-centred account underwritten by the inviolable status of those victimised), and a minimising/maximising account of agent-neutrality, underwritten by consequentialist considerations.

¹⁶ Notice that a crucial focus on the *author* relation simply equates the agent-neutral/relative distinction to the distinction between doing and happening. Agent-relative rules are explicitly understood as reasons to do or refrain from doing certain things *yourself*, i.e., act such that *you* don't murder people, and act such that *you* keep your promises, whereas agent-neutral rules are explicitly understood as reasons for things to happen, i.e., *you* must bring it about, to the greatest degree that you can, that there is minimal murdering in the world (even if this means murdering people yourself). This is an odd result, given that McNaughton and Rawling originally maintained that the doing/happening distinction was '*orthogonal*' to the neutral/relative distinction (McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction").

I have already highlighted the availability of non-consequentialist agent-neutral rules in Hammerton's logical space.¹⁷ In the next section I will argue that the notion of author agent-relativity that is said to underwrite deontological restrictions does not ground genuinely *deontological* restrictions. Rather, the substantive notion of author agent-relativity that is said to underwrite the existence of agent-relative deontological restrictions is most viably understood as grounding agent-relative *consequentialist* restrictions.¹⁸

5 Author Agent-Relative Restrictions are not Deontological Restrictions

At the end of section 1 I suggested that there was something odd about Hammerton's schematisation of deontological restrictions. To see this more clearly, consider:

(9) *Each agent must not murder people:*

(x) (x must act such that [x does not murder people])

As an exhortation directed to agents, (9) prohibits anyone and everyone from murdering people. Crucially, however, Hammerton maintains that each rule's content should be regarded, metaphorically, as what the rule is aimed at:

(10) *Each agent must bring about a world in which they do not murder people.*

(x) (x must bring about a world in which [x does not murder people])

But this is more than metaphorical. Rather, it encourages consequentialisation; the only way in which restrictions are said to generate different ultimate aims for different agents – *the sense in which the restriction is said to be agent-relative* – is to presume that restrictions direct different agents to bring about different possible worlds, i.e., *you* are required to bring about a world in which *you* don't murder people, and *I* am required to bring about a world in which *I* don't murder people. Moreover, if the deontic status of our actions is determined by a combination of the different worlds they instantiate along with the fact that agents have a special responsibility to bring about certain worlds themselves – *qua authors of their own actions* – then the most plausible rationale for grounding the agent-relativity of

¹⁷ Given Hammerton's distinction is meant to taxonomize regions of logical space, it certainly fails in this task (Dougherty, "Agent-Neutral Deontology", p. 531).

¹⁸ Incidentally, Hammerton has recently argued that deontological constraints (restrictions, special obligations, and permissions) *are*, in fact, maximising rules (not direct rules) and that 'a strong case can be made that consequentialism provides the best account of deontic constraints' (Matthew Hammerton, "Deontic Constraints are Maximizing Rules, *The Journal of Value Inquiry*, Vol. 54 (2020), pp. 571-588). I'll discuss the implications of this further in section 6. For the moment, I'll simply note that Hammerton's most recent position supports my thesis that the underlying rationale for agent-relative deontological restrictions is not, in fact, deontological at all. If deontological restrictions are to be *necessarily* agent-relative restrictions, then they are better characterised as agent-relative consequentialist restrictions, not deontological restrictions.

such restrictions will need to appeal to an axiological notion of *evaluator-relativity* to give significant content to the relativisation: the idea that each agent must bring about the world that is *best-relative-to-themselves* – *qua author of their own actions* – in a given circumstance.¹⁹

This is controversial, but, given that restrictions are said to be agent-relative in the sense that agents ought to have a special concern for what they *do* (the worlds they bring about), as opposed to what they merely allow to *happen* (the worlds they allow to occur), then there needs to be some significant content captured by the relativisation to the agent; the restrictions are author agent-relativised by maintaining that the state of affairs in which *you* murder someone is *worse-relative-to-yourself*, i.e., from your evaluator-relative perspective, than the world which you allow two (or more) people to be murdered by other people. Otherwise, the occurrence of ‘*x*’ within the square brackets bound by the initial quantifier is vacuous.

This misconceives the nature of deontological restrictions. To act such that you bring about a world in which you do not murder people is to maintain that the deontic status of your action is determined by the value of the world it instantiates, rather than a universal restriction on a specific action-type. The traditional deontological thought is simply that it is morally *wrong* for you to violate a restriction, not that your restriction violations are evaluator-relative bads.²⁰ Deontology has nothing to do with ultimate aims connected to the bringing about of possible worlds. Indeed, there is a perfectly good, non-metaphorical sense in which the deontological restriction ‘Do not murder’ gives anyone and everyone the same aim of not murdering and contains no back-reference to the agent for whom it is a reason.

6 A First Objection

There are two objections that can be raised at this point. The first concerns whether (9) can, in fact, be legitimately translated into (10), i.e., that there are counterexamples to the claim that (9) and (10) are equivalent, and that (10) does not require that an agent bring about a relevant outcome.

Consider the following:

¹⁹ For more on evaluator-relativity see Amartya Sen, “Evaluator Relativity and Consequentialist Evaluation”, *Philosophy & Public Affairs*, Vol. 12, No. 2 (1983), pp. 113-132; Douglas Portmore, *Commonsense Consequentialism* (Oxford: Oxford University Press, 2011); Douglas Portmore, “Agent-Centred Restrictions”, in Hugh LaFollette (ed.) *The International Encyclopedia of Ethics*, First Edition (Blackwell Publishing Ltd., 2013a).

²⁰ Granted, nothing in the formal distinction commits Hammerton to this proposal. However, as I noted in note 18, Hammerton has recently defended the view that deontic constraints *are* agent-relative consequentialist maximising rules (Hammerton, “Deontic Constraints are Maximizing Rules”). My view is that this is the most plausible rationale for capturing the agent-relativity of a restriction. Though, again, my view is that *deontological* restrictions are not agent-relative restrictions. The agent-relative consequentialist may well be committed to the view that restrictions are agent-relative consequentialist restrictions, and that’s fine (Portmore, *Commonsense Consequentialism*; Portmore, “Agent-Centered Restrictions”, p. 4). My suggestion is that the deontologist take restrictions to be fundamentally agent-neutral.

[S]uppose that Lee aims his rifle at the President's head and pulls the trigger with the intention and expectation that in doing so he will assassinate the President. Furthermore, suppose that, by chance, the President turns his head just as Lee fires, fortuitously moving it out of the path of the oncoming bullet and narrowly escaping the assassination attempt. In such a scenario, it is true that Lee has not done a certain act (the act of killing the President), yet it seems false that he has ensured that he did not do this act. After all, he has tried his best to make it the case that he did do this act, and it is odd to describe a circumstance where an agent earnestly tries to perform an act but fails due to bad luck as one where 'he ensures that he does not do the act.'²¹

The description of this event is accurate. Lee has not performed a prohibited action without ensuring (bringing it about that) he does not perform it. Nevertheless, this example is not sufficient to show that *adherence* to the deontological restriction on murdering people is not, by Hammerton's own admission, equivalent to the agent ensuring that they bring about a possible world (*with a view to ensuring the agent-relativity of the restriction*).²² Indeed, the restriction must be understood in this way if it is to be understood as agent-relative in the sense of preserving Parfit's 'intuitive idea' that agent-relative rules give each agent different ultimate aims – this is, after all, the intuition the Hammerton's formalisation is designed to capture.

To see this more clearly, recall the analytical equivalence of the consequentialist minimising-rules (7) and (8):

(7) (x) (x must bring it about, to the greatest degree that x can, that [there are minimal murders in the world])

(8) (x) (x must act such that [x minimises murders in the world])

Hammerton does not explain how to deal with such equivalences – he stipulates that reference to the agent within the square brackets of (8) is *eliminable*. Presumably, then, reference to the agent within the square brackets of (8) is eliminable because it is *superfluous*.²³ Why is it superfluous? Because, for all agents, x , *the world in which x must bring it about, to the greatest degree that x can, that there are minimal murders* is identical to *the world in which x must act such that x minimises murders in the world*. This is the only response available that enables one to hold on to the view that this type of consequentialist minimising-rule is genuinely agent-neutral – 'a rule is agent-neutral iff it is not agent-relative' – and it relies on understanding the ' *x must act such that*' locution as equivalent to the ' *x must bring it about*' locution. Likewise, in the case of restrictions, in order to understand the restriction on murdering people as *necessarily* agent-relative – in the sense that it

²¹ Hammerton, "Distinguishing Agent-Relativity from Agent-Neutrality", p. 243.

²² Lee has every *intention* of violating the restriction which is precisely why his act is subject to moral (and legal) condemnation.

²³ This is how Portmore deals with this kind of transformation (Douglas Portmore, "Agent-Relative vs. Agent-Neutral" in Hugh LaFollette (ed.) *The International Encyclopedia of Ethics*, First Edition (Blackwell Publishing Ltd., 2013b).

gives each agent a different ultimate aim of not murdering people themselves – *qua author of their own actions* – the ‘*x must act such that*’ locution must be treated as equivalent to the ‘*x must bring it about that*’ locution.

I am not saying that we *should* treat these as equivalents, but I am asserting that it is necessary to treat them as equivalents to ensure the relativisation of the restriction in terms of giving different agents different ultimate aims, and, due to concerns raised by a second objection that I am about to discuss, this is a high price to pay.

7 A Second Objection

A second objection concerns the idea that rules like (9) and (10) do not need to be explained by evaluator-relative values, i.e., philosophers can hold that *each agent must act such that she does not murder*, or that *each agent must bring about a world in which she does not murder* while rejecting an evaluator-relative consequentialist explanation of why this is true. She may, for instance, explain it via an appeal to the inviolability of moral persons (we must act this way because of their inviolability). Or, she might explain it via an appeal to agent-neutral teleology (life is intrinsically valuable, and the correct way to respond to this value is to always act in a way that honours the value of life where acting such that you do not murder people is a necessary feature of honouring the value of life).

I stressed that rules like (9) and (10) do not *need* to be explained by evaluator-relativity. The point is, rather, that Hammerton’s formalisation does not rule this out. In fact, it encourages such a reading, which is precisely what it is *not* meant to do (it is meant to demarcate a necessarily agent-relative deontology from an agent-neutral consequentialism). Again, if the deontic status of our actions is determined by a combination of the worlds they instantiate, along with the fact that agents have a special responsibility to bring about certain worlds themselves *qua author of their own actions*, then the underlying rationale for grounding the restrictions will have to appeal to an axiological notion of *evaluator-relativity*, i.e., the idea that each agent must act such that *they* bring about the world that is *best-relative-to-themselves* (from their perspective) in a given circumstance. Given that the restrictions are agent-relative in the sense that agents ought to have a special concern for what *they do* – *qua authors of their own actions* – as opposed to what they merely allow to *happen*, then evaluator-relativity explains how the restriction is relativised. If no appeal to evaluator-relative value is made, then the relativisation remains empty, or trivial. Indeed, deontological restrictions *already* prohibit the performance of certain token action-types *by specific agents* because those action-types are inherently wrong – who else is going to perform the token action-type in question?²⁴ To the extent that what matters is that *you* or *I* don’t perform the token action-type in question, the restriction is already relativised. The point is that Hammerton’s formalisation requires some additional content within the authorship parameter, and this is most naturally explained via the axiological notion of evaluator-relativity. Given the

²⁴ I’ll defend this point in more detail below.

trivial fact that the performance of a prohibited type of action by a particular agent will be the performance of that action-type by that particular agent, the fact that you and I must not do these things, *qua author of our own actions*, adds nothing to the equation.²⁵

If preserving the author agent-relativity of deontological restrictions comes at the potential cost of collapsing into agent-relative consequentialism, then so much the worse for the demarcation project. Alternatively, there is an intuitively plausible, pre-theoretical way of understanding the distinction between agent-neutral and agent-relative rules that is easily articulated in ordinary language, requires no formalisation, regiments no counterintuitive intuitions, and does everything a distinction between agent-relative and agent-neutral rules needs to do.²⁶

8 Enough with the Formalities?

The above discussion has revealed that Hammerton's formulation of the distinction between agent-neutral and agent-relative rules is designed to contrast two theoretically substantive intuitions: an author agent-relative account of deontology (which I've argued is not deontological at all), and an agent-neutral account of consequentialism tied to an instrumental and minimising/maximising account of practical rationality. Yet, on the face of it, a pre-theoretical distinction between agent-neutral and agent-relative rules is remarkably easy to articulate: namely, a distinction between universal moral rules which have a restricted scope of practical application (agent-relative rules), and universal moral rules which have a general, or wide scope of practical application (agent-neutral rules).

Within the context of the current discussion, there is a clear difference between the rules associated with deontological restrictions and the rules associated with deontological obligations and permissions. Deontological restrictions have a wide scope of practical application. They are rules that apply to anyone and everyone. No matter who you are (your numerical identity is irrelevant) or in what relations you stand to anyone or anything, you are prohibited from violating a deontological restriction. Not *because* you have a special responsibility not to violate deontological restrictions yourself (even though you do), but *because* of the inviolable status of the would-be victim(s). This is what makes deontological restrictions such as 'Do not murder', 'Do not tell lies', 'Do not steal', 'Do not cheat', etc., agent-neutral restrictions. Likewise, a consequentialist rule requiring anyone and everyone to maximise

²⁵ If anything, it is morally narcissistic to ground the force of deontological restrictions in the view that agents have a special responsibility not to violate said restrictions themselves. Certainly, each agent is responsible for not violating restrictions themselves, but to ground the force of the restriction in the fact that there is an allegedly 'ineliminable' occurrence of 'x' in some square brackets bound by an initial universal quantifier entails an odd view about that agent's moral importance, as well as positing as an inherent (and odd) connection between deontology and self-identity. See Alan Thomas, "Values, Reasons and Perspectives", *Proceedings of the Aristotelian Society*, New Series, Vol. 97 (1997), pp. 61-80.

²⁶ It was McNaughton and Rawling who insisted that 'a formalisation is only as good as the intuitions it regiments' (McNaughton and Rawling, "Agent-Relativity and Terminological Inexactitudes", p. 320).

general well-being will also be an agent-neutral rule requiring no particularising reference to any particular agent – no matter who you are, or what relations you stand to anyone or anything, the rule requires you to maximise general well-being.

Deontological obligations and permissions are different. While still universally applicable, these rules are only applicable to those agents who stand in the appropriate relation to an *object* of moral concern, such as a child, or a promise, or a personal project. After all, you are only obligated to keep *your* promises or look after *your* children insofar as *you* have made a promise or have children of *your own*, and you are only permitted to pursue your personal projects insofar as they are *your* projects. This is what makes deontological obligations and permissions such as ‘Look after *your* children’, ‘Keep *your* promises’, ‘Pursue *your* ambitions’ agent-relative rules.²⁷ Likewise, a rationally egoistic consequentialist rule requiring you to maximise your own well-being will also be agent-relative to the extent that it requires you to maximise your own well-being.

Incidentally, this proposal isn’t particularly novel. In line with Philip Pettit’s ‘full specification’ version of the distinction, we might say that agent-relative rules cannot be fully specified without *pronominal back-reference* to the person for whom it is a rule, where the motivating consideration (the rule) involves reference to some ‘object’ of the agent’s. Agent-neutral reasons, on the other hand, appeal to ‘certain properties of the circumstance, the action, the likely consequences of the action, or whatever’.²⁸

Similarly, in line with Nagel’s ‘general form’ version of the distinction, we might state that if a rule does not include an essential reference to the person for whom it is a rule, then it is an agent-neutral rule.²⁹ For example, ‘Reduce the amount of wretchedness in the world’. If, on the other hand, a rule does include an essential reference to the person for whom it is a rule, then it is an agent-relative rule. For example, ‘Do what makes *you* happy’.

I’ll explain how these ideas relate to the demarcation project in the next section, but, essentially, the distinction between agent-neutral and agent-relative rules concerns the practical scope of the rule in question and has nothing, inherently, to do with the distinction between deontological and consequentialist normative theories.³⁰

²⁷ Certainly, one can conceive of agent-relative restrictions, but they will inherit their agent-relativity from a pre-existing obligation or permission. You are prohibited from neglecting *your* children only insofar as they are *your* children and you’re obligated to care for them. Likewise, I can be prohibited from lying to *my* girlfriend only insofar as she is *my* girlfriend, and I am obligated to be truthful to her.

²⁸ Philip Pettit, “Universalizability without Utilitarianism”, *Mind*, New Series, Vol. 96, No. 381 (1987), pp. 74–82, p. 75. Pettit talks in terms of *reasons* (qua true propositions) not rules. I’ll say something more about this issue in note 31.

²⁹ Nagel, *The View from Nowhere*, pp. 152–153.

³⁰ Skorupski, *Ethical Explorations*, Ch. III; Buckland “Skorupski and Broome on the Agent-Neutral/Relative Distinction”. Perhaps the more interesting demarcation issue between the deontologist and the consequentialist is the deontologist’s rejection of the consequentialist’s teleological conception of practical reason (see Paul Hurley, “Consequentializing and Deontologizing: Clogging the Consequentialist Vacuum” in Mark Timmons (ed.) *Oxford Studies in Normative Ethics*, Vol. 3 (2013), pp. 123–153; Paul Hurley, “Comments on Douglas Portmore’s *Commonsense Consequentialism*”, *Philosophy and Phenomenological Research*, Vol. LXXXVIII, No. 1 (2014), pp. 225–232.

9 Historical Redux

Before wrapping up, I want to talk a little more about the implications of my proposal, what this means for Hammerton's distinction and McNaughton and Rawling's demarcation project, and what I mean by stating that the distinction between agent-neutral and agent-relative rules has nothing *inherently* to do with the distinction between deontological and consequentialist normative theories.³¹

I stressed in the introduction that McNaughton and Rawling's version of the distinction arose as an attempt to demarcate the boundaries of deontology and traditional universalist maximising act-consequentialism. 'Consequentialism', in McNaughton and Rawling's sense, provides an agent-neutral account of both the right and the good: the value of any state of affairs does not depend on the personal point of view of the agent, so 'no reference to the agent or her position in the world need enter into a consequentialist understanding of what makes an action right or

³¹ There are of course numerous ways of formalising the distinction between agent-neutral and agent-relative rules (or reasons) - nearly sixty years' worth of literature, in fact. Chronologically: E. A. Gellner, "Ethics and Logic", *Proceedings of the Aristotelian Society*, Vol. 55, No. 1 (1958), pp. 157-178; Thomas Nagel, *The Possibility of Altruism* (Oxford: Clarendon Press, 1970), Ch. X; Parfit, *Reasons and Persons*; Nagel, *The View from Nowhere*; McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction"; McNaughton and Rawling, "Honoring and Promoting Values"; McNaughton and Rawling, "Deontology and Agency"; Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell Books), Ch. 10, Ch. 11; McNaughton and Rawling, "Value and Agent-Relative Reasons"; John Skorupski, "Agent-Neutrality, Consequentialism, Utilitarianism: A Terminological Note", *Utilitas*, Vol. 7, No. 1 (1995), pp. 49-54; McNaughton and Rawling, "Agent-Relativity and Terminological Inexactitudes"; John Broome, "Skorupski on Agent-Neutrality", *Utilitas*, Vol. 7, No. 2 (1995), pp. 315-317; Skorupski, "Neutral Versus Relative: A Reply to Broome, and McNaughton and Rawling", *Utilitas*, Vol. 8, No. 2, pp. 235-248; B. C. Postow, "Agent-Neutral Reasons: Are They For Everyone?", *Utilitas*, Vol. 9, No. 2 (1997), pp. 249-257; Skorupski, *Ethical Explorations*, Ch. III; Douglas Portmore, "McNaughton and Rawling on the Agent-relative/Agent-neutral Distinction", *Utilitas*, Vol. 13, No. 3 (2001), pp. 350-356; McNaughton and Rawling, "Achievement, Welfare and Consequentialism"; McNaughton and Rawling, "Conditional and Conditioned Reasons", *Utilitas*, Vol. 14, No. 2 (2002), pp. 240-248; Mark Schroeder, "Reasons and Agent-Neutrality", *Philosophical Studies*, Vol. 135, No. 2 (2007), pp. 279-306; John Skorupski, *The Domain of Reasons* (Oxford: Oxford University Press, 2010), Ch. 3; Toni Rønnow-Rasmussen, *Personal Value* (Oxford: Oxford University Press, 2011), Ch. 9; John Broome, *Rationality Through Reasoning* (Oxford: Wiley Blackwell, 2013), Ch. 4; Portmore, "Agent-Relative vs. Agent-Neutral"; Dougherty, "Agent-Neutral Deontology"; Michael Ridge, "Reasons for Action: Agent-Neutral vs. Agent-Relative" in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*: (Fall 2017 Edition) URL = <https://plato.stanford.edu/archives/fall2017/entries/reasons-agent/>; Jamie Buckland, "Normative Reasons Qua Facts and the Agent-Neutral/Relative Dichotomy", *Philosophia*, Vol. 45 (2017), pp. 207-225; Krister Bykvist, "Agent-Relative and Agent-Neutral Reasons" in Daniel Star (ed.) *The Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press, 2018); Buckland, "Skorupski and Broome on the Agent-Neutral/Relative Distinction"; Hammerton, "Distinguishing Agent-Neutrality from Agent-Relativity". My sole concern has been with the formalisation *qua rules*, not reasons *qua* considerations (facts, true propositions, or obtaining states of affairs) with normative features counting in favour of action. Nevertheless, I do think the scope-based distinction I have drawn can be mapped to a suitable distinction drawn in terms of facts. For example, the fact that *your* child is hungry is an agent-relative reason for *you* to feed them, and the fact that *you* have promised to pay me for fixing your roof is an agent-relative reason for you to pay me, the fact that people are starving in distant lands is an agent-neutral reason for anyone give money to UNICEF (see Buckland, "Normative Reasons Qua Facts and the Agent-Neutral/Relative Dichotomy").

wrong'.³² A common criticism of agent-neutral accounts of the right and the good is their failure to accommodate the constraints that are central to deontological pluralism or commonsense morality. Agent-neutral consequentialism fails because it is too lax to capture the force of deontological restrictions and obligations, and too demanding or alienating in terms of the limits it places on an agent's deontological permissions.

Now, McNaughton and Rawling noted that consequentialism might be able to make room for threshold restrictions by claiming the performance of certain acts like murder and telling lies are intrinsically bad. The consequentialist can 'suck up' the intrinsic disvalue of killing innocents into their consequentialist vacuum cleaner, and, once the disvalue in question is in the dust bag or container (I presume a consequentialist would use a bagless vacuum cleaner in order to maximise suction), a threshold can be set as to where the intrinsic disvalue of killing innocents can be weighed against any further detrimental consequences of sticking to the initial restriction. However, McNaughton and Rawling argued that the vacuum cleaner interpretation fails to capture the *agent-relative* feature of deontological restrictions, i.e., the agent-neutrality that is central to consequentialist theories fails to account for *who* is breaching the restriction in the final calculation. For McNaughton and Rawling, then, deontological restrictions have a particularised concern with *the agent's* conduct: reference to the agent is an essential part of understanding why the action is wrong. This is precisely the author agent-relativity that (I have argued) leads us towards agent-relative consequentialism. Moreover, it is the *very same* rationale that McNaughton and Rawling extend to their analysis of deontological obligations and permissions to depart from maximising general well-being in order to, say, care for one's own children or pursue one's own personal projects. Indeed, what matters here, they argue, is not simply that looking after one's child is extremely valuable. Rather, what's important is that agents care for their own children *themselves*. It is this notion of author agent-relativity that McNaughton and Rawling insisted other leading formulations of the distinction (namely Nagel's 'general form' version and Pettit's 'full-specification' version) could not capture.³³

To understand this properly, consider the following rule:

Parents should look after their own children.

On my informal account, this rule is object agent-relative. It is a narrow scope rule telling each agent to look after their own children (presumably by *themselves* where possible – there is an extent to which constantly dumping your children on babysitters or grandparents instead of spending quality time with them *could* be construed as simply not caring for them in the appropriate way).³⁴ However, McNaughton and Rawling claim that this rule remains open to agent-neutral consequentialist interpretation. On an agent-neutral reading, the consequentialist will use

³² McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction", p. 168.

³³ McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction", pp. 170-175.

³⁴ I do not pretend for one moment to know anything about balancing the demands of work and parenting. But the fact that what constitutes caring for one's child *oneself* is such a complex issue lends support to the idea that author agent-relativity, in this context, is very difficult to generalise.

their vacuum cleaner to suck up the impersonal value of parents looking after their own children; if it is valuable for parents to look after their own children, then the agent-neutral rule requires parents to neglect their own child in order to maximise the total number of children looked after by their own parents. The agent-relative deontologist, on the other hand, maintains that the permissibility of neglecting one's child is determined by the fact *the parent* will be neglecting their child, not merely that a child is being neglected (notice, again, that the emphasis is on the author agent-relativity of the obligation not the object agent-relativity).

To cement this idea, they offer an example:

Consider, Alan, a conscientious subscriber to the doctrine that parents should look after their own children, who is reading a book to young Betty. In a full specification of his reason for reading to her it may be essential to mention that Betty is *his* daughter (which explains why he is reading to her and not some other child) and that *he* is doing the reading *himself* (which may explain why he is not employing someone else to do it). Neither of these pronominal back references to the agent is eliminable in a full account of his reasons for his reading to Betty. So Alan's reasons are agent-relative on this account. But nothing in our specification of Alan's beliefs determined whether Alan subscribes to the agent-relative or the agent-neutral version of this principle. So the full specification account fails to capture the distinction. [...] [Nagel's general form] account tries to avoid this problem by specifying that a reason is agent-neutral if it can be given a general form in which there is no reference to the person who has it. But what is the general form of Alan's reason? The obvious answer is that he is reading to Betty because he believes that parents should look after their children. This reason contains no reference to the agent and so is agent-neutral. But our difficulty was precisely that this principle could be given an agent-relative or agent-neutral reading.³⁵

It is unclear, however, what Pettit's or Nagel's (or my) distinction is failing to capture here.

McNaughton and Rawling are worried that we don't know whether Alan subscribes to an agent-relative or an agent-neutral reading of the rule, i.e., we don't know whether Alan is a consequentialist with a vacuum cleaner or a deontologist. They insist that Nagel tries to avoid this problem by specifying that a reason is agent-neutral if it can be given a general form that contains no reference to the person who has it, and stress that the obvious answer is that Alan is reading to Betty because he believes that parents should look after their own children – a reason they claim is agent-neutral because it contains no reference to the agent. But the general form of the rule '*Parents should look after their own children*' is clearly object agent-relative according to mine, Pettit's, and Nagel's formal account of the distinction in *The Possibility of Altruism* (where x ranges all agents and ϕ over all acts, events, and circumstances):

³⁵ McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction", pp. 171.

(11) (x, φ) (If φ is looking after x 's own child then x has reason to promote φ)

Granted, in *The Possibility of Altruism* Nagel maintains that agent-relative rules can be *subsumed* under their agent-neutral counterparts.³⁶ For instance:

(12) (x, φ) (If φ is looking after Alan's own child then x has reason to promote φ)

(13) (x, φ) (If $(\exists y)$ (φ is looking after y 's own child), then x has reason to promote φ)

But there is nothing about this that commits Alan to either of the readings; he can consistently subscribe to both. In his role as a parent, Alan feels it is his responsibility to read to Betty because he is a firm believer that parents should look after *their own* children. However, Alan may also feel there is something important about parents looking after their own children which should be promoted by *anyone* – even by those who don't have children themselves. Agent-relative rules assign certain *personal* or agent-relative values to things *for* individuals, i.e., the value to an individual of looking after *his* children. Yet, rules can also be construed agent-neutrally, but not in any consequentialist sense, but purely a sense in which they assign impersonal value to looking after one's own children. *Primarily*, then, the rules apply to conduct that directly supports the rule, such as reading a book to one's own child. Though, in assigning impersonal value to particular action types, the rules are also applicable *derivatively* to actions that promote similar conduct whether in oneself or from others. The impersonal value is conterminous to reasons of wide practical scope, i.e., the wide scope of the reason is explained by the impersonal value in the realisation of the state of affairs. Although, importantly, contra McNaughton and Rawling, even if the rule '*Parents should look after their own children*' is construed agent-neutrally, it does not require you or anyone else to neglect their own children in order to promote parents caring for their own children more generally. Nagel explicitly stresses this when referring to familial permissions:

Concern for one's wife and children is not merely concern for the welfare of some people whom one happens to be in a convenient position to help. Consequently, the [agent-neutralisation] of the reasons in whose acknowledgement this concern finds rational embodiment will not yield equal reasons to concern oneself with the wives and children of others. [...] The principle still yields [agent-neutral] rather than merely [agent-relative] reasons for people to look after their families. And even if one would defeat those reasons by directly assisting or forcing others to conform to them, two important consequences remain: there are reasons to avoid interference with others engaged in such

³⁶ Nagel originally drew a distinction between 'objective' and 'subjective' reasons, though later adopted Parfit's agent-neutral/relative terminology. I have replaced the terms in the quotations below to avoid confusion (see Nagel, *The Possibility of Altruism*, p. 90; Parfit, *Reasons and Persons*, p. 27, p. 143; Nagel, *The View from Nowhere*, p. 125).

activities, and also reasons to seek the social, economic, and political conditions which make pursuits possible, not only for oneself but for others.³⁷

Contra the demarcation project, then, there is nothing *inherently* consequentialist, utilitarian, or maximising about agent-neutral rules. There is, for example, nothing counterintuitive about a rule that provides agent-neutral reasons for anyone and everyone to minimise acts which involve the harming of others in order to benefit others. Indeed, Nagel was always keen to insist that ‘the principle of [agent-neutrality] does *not* automatically yield a species of utilitarianism, or some other counter-intuitive 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 individual human lives’.³⁸

10 But Doesn’t Nagel Think that Restrictions are Agent-Relative?³⁹

An immediate concern here is that, in *The View from Nowhere*, Nagel takes deontological restrictions to *be* agent-relative reasons (or rules).⁴⁰

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.⁴¹

That said, Nagel also acknowledged that the agent-relative status of restrictions is perplexing:

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.⁴²

³⁷ Nagel, *The Possibility of Altruism*, p. 130.

³⁸ Nagel, *The Possibility of Altruism*, p. 142.

³⁹ I have discussed these ideas elsewhere in Buckland “Skorupski and Broome on the Agent-Neutral/Relative Distinction”.

⁴⁰ Nagel refers to what I have been calling restrictions as constraints. Following Portmore, I take restrictions to be a subclass of the broader category of constraints encompassing restrictions, obligations, and permissions.

⁴¹ Nagel, *The View from Nowhere*, p. 175.

⁴² Nagel, *The View from Nowhere*, p. 175.

Nagel gets around this issue by appealing to the doctrine of double effect: the idea that there is a relevant distinction between *intentionally* bringing about an outcome and causing an outcome that is *foreseeable*, but not intentional:

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.⁴³

The rationale for a deontological restriction, then, is that its violation would be an instance of *intentional victimization* on the effected agent, but this is not an appeal to author agent-relativity in the sense that each agent is required not to commit intentional violations *herself—qua author of her own actions*:

[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.⁴⁴

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.⁴⁵ Moreover, it seems that Nagel actually abandoned the idea that the rationale for deontic restrictions could be sought in the evil intentions of the victimiser in favour of a rationale grounded solely in the inviolable status of those victimised:

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.⁴⁶

As far as Nagel's qualified position is concerned, then, the underlying rationale for deontological restrictions is agent-neutral: anyone and everyone is prohibited from, say, murdering *because* the potential victim possesses an inviolable right to life which is universally, but non-consequentially valuable (not *because* each agent

⁴³ Nagel, *The View from Nowhere*, p. 179. The doctrine of double effect encompasses a different agency rationale than the author agent-relative account of restrictions. The central feature of the author agent-relative account of restrictions is that the requirement that one not violate a deontological restriction is more stringent than a demand to prevent violations. The significance of agency associated with the doctrine of double effect, on the other hand, is that it is impermissible for an agent to *intentionally* violate a restriction in order to secure some greater good (see Richard Brooke, "Agency and Morality", *The Journal of Philosophy*, Vol. 88, No. 4 (1991), pp. 190-212.

⁴⁴ Nagel, *The View from Nowhere*, p. 184.

⁴⁵ See Eric Mack "Deontic Restrictions Are Not Agent-Relative Restrictions", *Social Philosophy & Policy*, Vol. 15, No. 2 (1998), pp. 61-83.

⁴⁶ Thomas Nagel, "Personal Rights and Public Space", *Philosophy & Public Affairs*, Vol. 24, No. 2 (1995), pp. 83-107.

has a special responsibility not to violate restrictions *qua authors of their own actions*).⁴⁷

11 A Final Objection

I have argued that the rules expressing deontological obligations and permissions are object agent-relative in an interesting sense, whereas deontological restrictions are agent-neutral restrictions; the author agent-relativity which is said to underwrite deontological restrictions is either trivial or encourages agent-relative consequentialisation. A final objection I want to consider is that there are, in fact, two kinds of agent-relativity (object and author), but that they share *enough* in common to both warrant the label ‘agent-relativity’.

Consider, again, the idea that a deontologist typically endorses an agent-relative restriction on murdering people:

(9) *Each agent must not murder people:*

(x) (x must act such that [x does not murder people])

An agent-neutral consequentialist contends that the only justification for such a rule is that killing is *bad*. Hence, (9) only holds because of:

(7) *Each agent must minimize murders:*

(x) (x must bring it about, to the greatest degree that x can, that [there are minimal murders in the world])

Or, when it comes to, say, parental caregiving, a traditional deontologist might contend that:

(13) *Each agent must care for her children:*

(x) (x must act such that [(y) (y is x’s child $\rightarrow$ x cares for y)])

An agent-neutral consequentialist contends that (11) only holds because parental caregiving is *good*:

(14) *Each agent must ensure that children are cared for by their parents:*

(x) (x must ensure that [(y) (z) (y is a child of z $\rightarrow$ z cares for y)])

When there is a conflict between (9) and (7) or (13) and (14), the agent-neutral consequentialist holds that the agent-neutral rule takes precedence, whereas the deontologist holds that, while we should discourage murdering and encourage

⁴⁷ Attempts to preserve the agent-relativity of restrictions are perhaps attempts to avert 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 the same type by others (or oneself). But this idea is only paradoxical if the advocate of deontic restrictions accepts that the disvalue or objectionable property of their violation is to be minimised in the first place.

parental caregiving, we should not do so at the expense of violating (9) or (13) *ourselves*.

Now, what makes (9) and (13) both ‘agent-relative’ (in a significant sense) is the fact that the *doer* is also the agent who is subject to the prescription, or the person to whom someone (or something) is specially related, or both – the person who should not murder, or who should care for her children. In the agent-neutral consequentialist’s directives (7) and (14), on the other hand, the *ensurer’s* aim should be that no one murders, and that all parents care for their children, i.e., the ensurer plays no special role in what she should aim to ensure. Hence, object and author agent-relativity do share ‘agent-relativity’ in an important sense, i.e., ‘*doer-relativity*’.

This, is, however, precisely the point of contention. As I’ve stressed, as far as deontological restrictions are concerned, they *already* prohibit the performance of specific action-types open to agents. Moreover, as I pointed out in note 7, this idea collapses the distinction between agent-neutral and agent-relative rules into the doing/happening distinction.

Recall, (again):

(7) (x) (x must bring it about, to the greatest degree that x can, that [there are minimal murders in the world])

(8) (x) (x must act such that [x minimises murders in the world])

I have argued that the analytical equivalence of these rules means that Hammerton’s distinction cannot remain structural: it must appeal to an agent-focused account of deontology underwritten by an account of author agent-relativity that encourages the consequentialisation of deontological restrictions.

Finally, there is the added danger that Hammerton’s distinction actually gets the wrong results, i.e., that the consequentialist rule directing agents to maximise general well-being is best understood as an agent-relative rule. Hammerton needs to be in a position to give priority to (7) over (8). However, as Jonathan Dancy noted in an early critique of McNaughton and Rawling’s distinction, there is nothing in Hammerton’s distinction that allows him to do so. Following Dancy, we can utilise an asymmetry criterion which permits us to maintain that the agent-relative rule holds *because* the agent-neutral rule holds, i.e., we can ask ourselves if there is a sense in which one rule *leads* and the other *follows*. If (8) holds because of (7), then we could proclaim the rule as agent-neutral. However, it seems as if (7) holds *because* (8) does. As a rule directed at individual agents, (7) directs each agent to perform the action with the ‘best’ outcome available to *them* – the rule is (metaphorically) *aimed at* maximising general well-being. Because of this, if another agent can do better than you, then you should let them. Or, if you can arbitrate between two agents, then you should prefer the one whose action will be ‘better’. However, you operate in this manner *because* it is required of you by (8).⁴⁸

⁴⁸ Dancy, *Moral Reasons*, p. 205.

12 Conclusion

I have argued that attempts to demarcate the boundaries of traditional direct act-consequentialism and deontology via an appeal to a formal distinction between agent-neutral and agent-relative rules fails insofar as agent-relative deontological restrictions are best not thought of as deontological restrictions at all. The formal distinction on which the success of the demarcation projects depends leads to the conclusion that the purportedly agent-relative deontological restrictions are, in fact, best thought of as agent-relative consequentialist restrictions.

Furthermore, I have defended both the possibility as well as the plausibility of agent-neutral deontological restrictions. These restrictions are agent-neutral restrictions in the sense that they apply to anyone and everyone regardless of who they happen to be, or the relations in which they stand to anyone or anything. They are deontological restrictions in the sense that they universally prohibit the performance of specific action-types by anyone and everyone because the performance of those action-types is inherently wrong. Contra the orthodox view, then, deontological restrictions are not, necessarily, agent-relative restrictions, and neither is deontology. In fact, given that most plausible deontological normative theories will consist of various agent-neutral restrictions, agent-relative obligations, and agent-relative permissions, it is questionable whether ‘deontology’ qua normative theory can be suitably characterised as exclusively agent-neutral either.

Acknowledgements I am extremely grateful to an anonymous referee for their helpful comments, criticisms, and suggestions. I am also very grateful to Brad Hooker, Chris Jay, Barry Lee, Mary Leng, Christian Piller, Alan Thomas, Wlodek Rabinowicz, Philip Stratton-Lake, and audiences of the practical philosophy group at the University of York for comments on presentations from which this paper was derived.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

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