

Patient-Relativity in Morality*

Matthew Hammerton

It is common to distinguish moral rules, reasons, or values that are agent-relative from those that are agent-neutral. One can also distinguish moral rules, reasons, or values that are moment-relative from those that are moment-neutral. In this article, I introduce a third distinction that stands alongside these two distinctions—the distinction between moral rules, reasons, or values that are patient-relative and those that are patient-neutral. I then show how patient-relativity plays an important role in several moral theories, gives us a better understanding of agent-relativity and moment-relativity, and provides a novel objection to Derek Parfit’s “appeal to full relativity” argument.

INTRODUCTION

The distinction between an agent-relative reason, rule, or value, and an agent-neutral reason, rule, or value is widely recognized as one of the most important distinctions in value theory. A second related distinction is that between a moment-relative reason, rule, or value and a moment-neutral reason, rule, or value. Although this distinction is not as prominent in value theory as the agent-neutral/relative distinction, it is also important and has been discussed by several philosophers. In this article, I introduce a third distinction that stands alongside these two distinctions—the distinction between a patient-relative reason, rule, or value and a patient-neutral reason, rule, or value. This distinction is based on the observation that the numerical identity of a moral patient may be a morally

* I would like to thank Daniel Nolan and Nic Southwood for invaluable advice on this article. I am also very thankful for helpful comments from Ryan Cox, Philip Pettit, Doug Portmore, Daniel Stoljar, and Kalpana Vignehsa, as well as two anonymous referees and the associate editors of this journal.

Ethics 127 (October 2016): 6–26

© 2016 by The University of Chicago. All rights reserved. 0014-1704/2016/12701-0002\$10.00

significant factor. In some contexts the moral significance of patient identity has already been noted. However, a sustained and systematic account of this phenomenon has not been provided. I offer such an account in this article. Furthermore, I show how this account leads to a distinction between patient-relativity and patient-neutrality. Recognizing the possibility of patient-relativity is important because it plays a key role in several moral theories, gives us a better understanding of agent-relativity and moment-relativity, and provides a novel objection to Derek Parfit's "appeal to full relativity" argument.

Here is how the rest of this article will proceed. In Section I, I will explain the agent-neutral/agent-relative distinction, and in Section II I will explain the moment-neutral/moment-relative distinction. This provides a helpful background for the discussion that follows in Section III. In part *A* of Section III, I introduce the patient-relative/patient-neutral distinction by presenting several cases where a patient's numerical identity appears to make a difference to our moral obligations. I follow this in part *B* with a discussion of the logical structure of patient-relativity and its relation to agent-relativity and moment-relativity. Finally in part *C* I show how several moral theories may be committed to patient-relative rules. I then show how the possibility of patient-relativity provides a response to Derek Parfit's "appeal to full relativity" argument.

I. AGENT-RELATIVITY AND AGENT-NEUTRALITY

The agent-neutral/agent-relative distinction was brought to the attention of contemporary philosophers by Thomas Nagel's discussion of it in *The Possibility of Altruism*, although it was not until Parfit's *Reasons and Persons* that the distinction came to be described using the terminology "agent-neutral" and "agent-relative."¹ Since Nagel's initial discussion, many different accounts of the distinction have been put forward.² Debates about how to formulate the distinction are complicated by the fact that it can be applied to three different categories—reasons, rules, and values. Which of these categories, if any, should take priority in explicating the distinction is often disputed. Importantly, the participants in these debates agree that they are talking about the same distinction and give roughly the same assessment of various cases as agent-neutral or agent-relative. What they disagree about is the optimal way of cashing the distinction out.

1. Thomas Nagel, *The Possibility of Altruism* (Princeton, NJ: Princeton University Press, 1970); Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984).

2. For detailed discussion of the many ways of formulating the distinction, see Michael Ridge, "Reasons for Action: Agent-Neutral vs. Agent-Relative," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Winter 2011 ed.), <http://plato.stanford.edu/archives/win2011/entries/reasons-agent/>; and Douglas Portmore, "Agent-Relative vs. Agent-Neutral," in *International Encyclopedia of Ethics*, ed. Hugh LaFollette (London: Wiley Blackwell, 2014).

Which of the many formulations of the distinction is best is not a question I will address in this article. However, I will review two popular accounts of the distinction as it applies to moral rules as a set-up to the discussion that follows. It should be noted that in choosing formulations that apply to moral rules I am not taking a stand on whether rules have priority over reasons or values. In this article, every claim I make about the neutrality or relativity of moral rules can be translated into a claim about the neutrality or relativity of reasons or values.

According to Derek Parfit, an agent-relative rule is a rule that gives different agents different aims, whereas an agent-neutral rule is a rule that gives all agents the same aim.³ As an example, consider a moral rule that prohibits the killing of innocent people which I will call “DK,” short for “Do not kill”:

DK: Each agent must not kill innocent people.

DK is an agent-relative rule because it gives each agent a different aim—the aim that he or she does not kill. Thus, DK gives Alfred the aim that “Alfred does not kill” and Sita the aim that “Sita does not kill.” DK can be usefully contrasted with a rule that instructs us to minimize the killing of innocent people. Let’s call this rule “MK,” short for “minimize killing”:

MK: Each agent must minimize the killing of innocent people.

MK is agent-neutral because it gives each agent the same aim. Alfred’s aim is that there is as little overall killing as possible, and Sita’s aim is the same.⁴ This contrast is most vivid in cases where, by killing an innocent

3. Parfit, *Reasons and Persons*, 27, 129.

4. Two clarifications are necessary here. First, when Parfit talks of agents having the same aim, he is referring to “ultimate aims” and not “derived aims.” If Alfred and Sita are in different circumstances, then they might find that minimizing killing requires them each to do different things. As such, MK gives them different aims. However, these different aims are derived from the same ultimate aim that they each share—the aim of minimizing killing. Second, one might interpret MK as giving different agents different ultimate aims. Perhaps it gives Alfred the ultimate aim that “Alfred minimizes killing” and Sita the ultimate aim that “Sita minimizes killing.” However, as David McNaughton and Piers Rawling, “Agent-Relativity and the Doing-Happening Distinction,” *Philosophical Studies*, 63 (1991): 167–85 note, the distinction can still be made even if such an interpretation of MK is legitimate. For, although this interpretation counts MK as giving different agents different ultimate aims, there is also an equally valid interpretation of MK where it gives different agents the same aim (i.e., the aim of minimizing killing). However, there is no plausible interpretation of DK where different agents have the same ultimate aim. Thus, Parfit’s account of the distinction is best stated as follows: A rule is agent-neutral if and only if there is a valid interpretation of it that gives all agents the same ultimate aim, and is agent-relative otherwise.

person, an agent can prevent more killings by others. Suppose the Mafia make a credible threat to Alfred that if he does not kill a random innocent person they will kill several innocent people. What should Alfred do? According to DK, Alfred must not kill the innocent person even though that will mean there is more killing overall. This is because DK gives Alfred a special concern with his own killing, requiring him to give it priority over the killings that others might commit. According to MK, Alfred must kill the innocent person. For, in this situation, doing so will result in fewer killings overall and MK requires Alfred to give equal consideration to any killings that may occur whether or not he is responsible for them.

Parfit's intuitive account of the distinction can be contrasted with the more formal account of David McNaughton and Piers Rawling. According to McNaughton and Rawling, moral rules like DK and MK have an underlying logical form that is not immediately apparent.⁵ For example, they interpret DK as having the following underlying structure:

- (1) $(x)(x \text{ should ensure } [(x \text{ does not kill})])$.

Whereas MK has this structure:

- (2) $(x)(x \text{ should ensure } [(y)(y \text{ does not kill})])$.

McNaughton and Rawling argue that (1) is agent-relative because the content of the rule (the sentence in the square brackets) makes reference to a specific agent. By contrast, (2) is agent-neutral because the content of the rule does not refer to any specific agent. More formally, they say that a rule expressed with the form $(x)(x \text{ should ensure } [. . .])$ "is agent-relative if and only if there is an occurrence of x in the square brackets bound by the initial universal quantifier; it is agent-neutral otherwise."⁶

Whatever account of the agent-neutral/agent-relative distinction one accepts, there appears to be widespread agreement on its importance. Thomas Hurka describes it as one of "the greatest contributions of recent ethics."⁷ Several others have claimed that it gives a precise and illuminating account of the distinction between consequentialism and nonconsequentialism.⁸ Finally, James Dreier has argued that it is distinct

5. McNaughton and Rawling, "Agent-Relativity and the Doing-Happening Distinction."

6. David McNaughton and Piers Rawling, "Honoring and Promoting Values," *Ethics* 102 (1992): 835–43, 837.

7. Thomas Hurka, "Moore in the Middle," *Ethics* 113 (2003): 599–628, 628.

8. For example, see McNaughton and Rawling, "Honoring and Promoting Values"; and Philip Pettit, "The Consequentialist Perspective," in *Three Methods of Ethics*, by Marcia Baron, Philip Pettit, and Michael Slote (Oxford: Blackwell, 1997).

from, yet more fundamental to moral theory than the distinction between consequentialism and nonconsequentialism.⁹

II. MOMENT-RELATIVITY AND MOMENT-NEUTRALITY

The next distinction goes by many different names. John Broome uses the terms “moment-neutral” and “moment-relative” to mark it. Michael Ridge employs the terms “diachronic” and “synchronic.” Jennie Louise prefers the terms “temporal-neutral” and “temporal-relative.” Finally, Francis Kamm describes this kind of relativity as the “agent-centered and time-slice approach to persons” and “deontology of the present moment.”¹⁰ I will use Broome’s terminology of “moment-neutral” and “moment-relative,” although nothing important hangs on the terminology used. The best way to see this distinction clearly is to consider an alternative version of the Mafia case described above. Suppose that Alfred is not in a situation where the Mafia will kill several if he fails to kill one. Instead, he is in a situation where it is his own future self who will kill several if his current self fails to kill one. Perhaps Alfred finds his mind slowly being overtaken by homicidal desires and correctly judges that, unless he kills one now, his weak will and morally corrupted desires will soon lead him to kill several. According to the agent-relative rule DK, Alfred must not kill innocent people. However, in this scenario where he has to choose between killing one now and killing several later what should he do? One possible moral rule (let’s call it DK′) requires him to give equal consideration to all of his potential killings and prevent as many of them as possible. This would require Alfred to kill the one now to prevent his later killing of several. Such a rule is moment-neutral because it does not give priority to any particular moment in time such as the present. An alternative moral rule (let’s call it DK′′) requires Alfred to give special priority to any killing he is about to perform, and thus to not kill the one now. Such a rule is moment-relative because it gives special priority to the present moment. We can also state the difference between these two rules using Parfit’s talk of “aims” (although Parfit himself never does this). DK′ is moment-neutral because it gives different time slices of the same agent the same aim (the aim that the agent kills as little as possible). DK′′ is moment-

9. See James Dreier, “The Structure of Normative Theories,” *Monist* 76 (1993): 22–40, 39.

10. John Broome, *Weighing Goods* (New York: Blackwell, 1991); Michael Ridge, “Agent-Neutral Consequentialism from the Inside-Out: Concern for Integrity without Self-Indulgence,” *Utilitas* 13 (2001): 236–54; Jennie Louise, “Relativity of Value and the Consequentialist Umbrella,” *Philosophical Quarterly* 54 (2004): 518–36; Francis Kamm, “Review: Non-consequentialism, the Person as an End-in-Itself, and the Significance of Status,” *Philosophy and Public Affairs* 21 (1992): 354–89.

relative because it gives different time slices of the same agent different aims (the aim that the time slice does not kill at the moment it inhabits).

In the example described above the moment-relative rule was also an agent-relative rule (the rule both gives different agents different aims, and different time slices of the same agent different aims). However, moment-relative rules can also be agent-neutral. For example, John Broome describes a possible moral rule that appears to be agent-neutral yet moment-relative.¹¹ According to this rule when a miner is trapped in a mine everyone is required to perform actions (e.g., contributing resources to the rescue effort) that will maximize the miner's chance of being saved. Stated thus the rule appears agent-neutral (it gives all agents the same aim). However, Broome's rule also has a caveat requiring agents to perform any action which contributes to saving miners who are currently trapped, even if doing so will mean that more trapped miners in the future will go unsaved. Thus, we must not hold resources back from current efforts even if doing so would allow more miners to be saved in the future. This caveat does not alter the rule's agent-neutral status (everyone has the same aim at any one time) but does appear to make the rule moment-relative (our different time slices are given different aims).

So it appears that the moment-neutral/moment-relative distinction is logically independent of the agent-neutral/agent-relative distinction. There are four possible ways that these two distinctions can combine.¹² First, a rule might be both agent-neutral and moment-neutral (e.g., the "maximize total net pleasure" rule of classical utilitarianism). Second, a rule might be both agent-relative and moment-relative (e.g., the traditional deontological constraint on killing). Third, a rule might be agent-neutral but moment-relative (e.g., the rule specified in Broome's miner example). Fourth, a rule might be agent-relative yet moment-neutral (e.g., a rule that normally prohibits killing to prevent more killing but permits it when current and future killers are the same agent).

Although the moment-neutral/moment-relative distinction has not received as much attention as the agent-neutral/agent-relative distinction, it is nonetheless an important distinction in value theory. For example, Francis Kamm uses moment-relativity to argue that deontological constraints are best explained by the inviolability of victims. Douglas Portmore uses it to argue that agent-relative consequentialism does better than its deontological rivals. Finally, John Broome argues that moment-relativity is more important than agent-relativity when it comes to understanding deontological side-constraints.¹³

11. Broome, *Weighing Goods*, 8.

12. See Louise, "Relativity of Value," 535.

13. See Kamm, "Non-consequentialism"; Douglas Portmore, *Commonsense Consequentialism* (New York: Oxford University Press, 2011), 103–8; and Broome, *Weighing Goods*, 9.

III. PATIENT-RELATIVITY AND PATIENT-NEUTRALITY

The two distinctions I discuss in Sections I and II are logically independent and yet appear to be closely related. Both distinctions are recognized by value theorists and applied to important debates in value theory. However, there is a third distinction that appears to belong to this family of distinctions and yet has so far gone unrecognized. I will label this distinction the “patient-neutrality/patient-relativity distinction.” The idea behind this distinction is simple. The existence of agent-relative moral theories shows us that the numerical identity of the agent acting might make a difference to the deontic status of the act. The existence of moment-relative moral theories shows us that the numerical identity of the time slice acting might make a difference to the deontic status of the act. A third factor that can be variable in such cases is the identity of a moral patient. In many scenarios when an agent faces a morally significant choice there is a moral patient who is affected in some way by the choice the agent makes. Perhaps, just as the numerical identity of the agent or the time slice can make a difference, the numerical identity of the moral patient can make a difference. When patient identity makes a difference to the deontic status of an act then we have a case of patient-relativity. When it makes no difference we have patient-neutrality.

This is a reasonable first attempt at explaining patient-relativity. However, much more needs to be said to explain this distinction and demonstrate its importance to moral theory. I will begin this task in part *A* below by presenting a case which, according to our commonsense moral judgments, appears to be an instance of patient-relativity. In part *B*, I will offer a precise account of the patient-relative/patient-neutral distinction. I will also look at different ways that patient-relativity can combine with agent-relativity and moment-relativity. Finally, in part *C*, I will show some useful applications in moral theory for patient-relativity.

A. An Initial Case of Patient-Relativity

Some moral theories endorse an agent-relative rule prohibiting promise breaking. Such a rule requires agents not to break a promise even if doing so is the only way to prevent more promise breaking by others. However, for any theory that endorses such a rule there is a further question we can ask. Should the rule be interpreted as moment-neutral or moment-relative? To answer this we ask whether an agent is permitted to break a promise now if that is the only way she can ensure that she keeps several equally weighty promises in the future. For example, suppose Alfred has made a promise to his colleague Marvin to send a report to Marvin by tonight. However, he has also made five separate promises to five other colleagues that require him to complete various tasks by tomorrow (which are each as important as sending Marvin the report). Alfred has failed to

make a start on any of these tasks and realizes that, unless he works on these tasks now, thereby breaking the promise to Marvin, he will break all five of tomorrow's promises. Thus he must choose between keeping his promise to Marvin now and breaking promises to five other colleagues tomorrow, or breaking his promise to Marvin now and keeping his promises to five others tomorrow. In these kinds of cases I suspect that some who endorse the agent-relative rule prohibiting promise breaking will favor a moment-neutral interpretation while others will favor a moment-relative interpretation. Those who favor moment-neutrality might reason as follows:

Those who make incompatible promises have done wrong and owe compensation to any promisee who fails to get her due. However, in deciding which promises to keep and which to sacrifice, it is best that the agent minimizes her overall promise breaking because, other things being equal, it is better that a lesser number rather than a greater number are wronged by the bad faith of promise breaking. Furthermore, temporal order does not seem to be a morally significant factor. Why should the fact that the promise made to you will be fulfilled before the promise made to me mean that the fulfilment of your promise counts for more than the fulfilment of mine? But, if temporal order is not morally significant, then it is morally unjustified to favor keeping the promise that will be broken first simply because of its temporal order.

By contrast, those who favor the moment-relative account may reason as follows:

The person who makes incompatible promises still has an obligation, at each moment, to keep any promise that it is in her power to keep. To fail to do so for the gain of more personal promise keeping overall is to use one of the promisees and the commitment she is owed as a means to the end of fulfilling other promises, and this is morally objectionable. Furthermore, my moment-neutral rival is already committed, via her agent-relativity, to the claim that maximizing general promise keeping is not a good enough reason for breaking a promise. But if she is committed to this then how can she say that maximizing personal promise keeping justifies breaking a promise? To think that the general case does not count but the personal case does is to be overly partial to oneself, and such partiality is morally objectionable.

Perhaps many will find the moment-relative account the more convincing of the two in this debate. However, consider a variation of the case where the one proximate promise and the five equally weighty future promises all have the same promisee (i.e., the same moral patient). For example, perhaps Alfred's one promise due tonight and five promises

due tomorrow are all promises he made to Marvin. Does the promisee being the same person rather than different people change things? Many will be inclined to say that it does. For, it is one thing to say that I must not break a promise to you now in order to keep (equally weighty) future promises I have made to five others. It is another thing to say that I must not break a promise to you now if that is the only means available to me to keep five (equally weighty) future promises I also made to you. In the former case you will be worse off if I break my promise. However, in the latter case you will be better off overall. For, from your perspective as the promisee, five promises kept and one broken is preferable to one kept and five broken (assuming that all promises are equally weighty). Furthermore, in the former case you might reasonably claim that you have been treated unfairly by having the promise owed to you neglected so that others can gain some advantage. However, in the latter case such claims of unfairness fail because the promise owed to you now is being neglected not to advantage others at your expense, but rather to your advantage. On these grounds, many who endorse the moment-relative interpretation of the constraint on promise breaking will agree that the constraint is best given a moment-neutral interpretation when we are making a trade-off between present and future promise breaking involving the same patient. Thus, in this case it appears that the identity of the moral patient makes a difference to moral requirements. When the patients involved in a trade-off are numerically identical then the constraint is moment-neutral. However, when the patients are not numerically identical then the constraint is moment-relative.

The example above involves promising. However, other moral duties also appear to have this patient-relative structure. For example, consider the special duties we owe to friends. Such duties are agent-relative because they require us not to neglect our friends even if doing so will cause less neglecting of friends by others. They also appear moment-relative because they require us not to neglect a friend now even if doing so will result in our neglecting other friends less in the long run. However, what about a case where the short-term and long-term trade-offs involve the same friend? Suppose that Sita has the options of either giving support to her friend Ruth now or withholding support at this moment. However, Sita learns that if she gives her support now then this will make her unable to give her support on many future occasions (perhaps giving support at this moment will lead to a deep resentment which will make it hard for her to support Ruth in the future). On the other hand, if Sita withholds support now she will be able to offer more support to Ruth in the future and will end up being a better friend in the long run. Many will think that, although duties of friendship are generally moment-relative, in cases like this where short-term and long-term trade-offs involve the same friendship (i.e., the same moral patient), the moment-relative con-

straint no longer applies. Thus, duties of friendship appear plausible candidates for patient-relativity.

The examples considered so far (promissory duties and duties of friendship) are both special duties. However, some natural duties are also candidates for patient-relativity. For example, consider a natural duty prohibiting stealing. Such a duty appears agent-relative because it requires each agent not to commit an act of theft even when doing so is the only way to prevent several comparable thefts being committed by others. Furthermore, many will hold that this duty is generally moment-relative because we must not steal from one person at the present moment even if doing so is the only way to prevent our future selves from stealing from several others. However, consider a case where the current and future thefts are to be committed against the same person. It is not easy to construct a plausible version of such a case, but here is an attempt. Suppose that Rudi has set up a scam to steal one hundred dollars from Suki's bank account on five separate occasions over the next month. The scam is difficult to undo and the only way for Rudi to prevent it from going ahead is by stealing one hundred dollars from Suki's account now (perhaps doing so will set off a warning system that will prevent the future thefts). In such a case many will claim that the agent should steal from the patient now to prevent several comparable thefts against that same patient in the future. However, they will insist that the agent must not steal from one patient now in order to prevent several comparable thefts against different patients in the future. Thus, duties not to steal property are another candidate for patient-relativity.¹⁴

At this point it is worth clarifying what these initial examples of patient-relativity are intended to show. I am not arguing that any of the patient-relative rules considered above are in fact true. I am merely pointing to the possibility of such rules to highlight patient-relativity as a possible feature that a moral theory can possess. Many will find patient-relative rules of the kind described above intuitively plausible and even compelling. However, others will reject them. Agent-neutral consequentialists will reject them because they reject all agent-relative constraints. Others will reject them because they prefer the moment-neutral interpretation of such constraints (even when the trade-off is between different patients). Finally, someone who endorses the moment-relative interpretation of such a constraint might insist that this interpretation holds even in same-patient cases. For example, it might be claimed that it is

14. Reflection on these examples raises an interesting question about whether there is a general principle that predicts patient-relativity (I thank an anonymous referee for suggesting this possibility). I will not propose such a principle in this article. However, in part *C* I will show how some moral theories are committed to patient-relativity, and this discussion will illuminate several directions we might take in formulating a general principle.

wrong to break a promise to one patient in the present moment, even if doing so is the only way to prevent a greater amount of future promise breaking to that same patient. For perhaps doing so is to disrespect the institution of promising and this is wrong regardless of whether it is harmful or unfair to any particular individual. It is not my concern here to offer replies to each of these critics. However, part *C* will indirectly address the issue of whether we should endorse any patient-relative moral rules by showing how some normative theories might entail, or provide strong support for, patient-relativity.

B. The Logical Structure of Patient-Relativity

The cases we considered above give us an initial picture of what a patient-relative moral rule might look like. However, more work needs to be done to get a clearer picture of the nature and scope of patient-relativity. In the cases we have already considered, the identity of the patient acted as a switch that turns on and off the moment-relative requirement. When the proximate patient was also the future patient affected by the action, then moment-relativity was turned off. When the proximate patient was distinct from any future patients, then moment-relativity was turned on. However, the same process may also occur in relation to agent-relativity. To see how consider the following example. Suppose that if I break a promise to you I will prevent several promises made by other agents to other patients from being broken. The agent-relative rule will prohibit promise breaking in such circumstances. However, let's modify the example and suppose that the promises made by other agents are also promises made to you. I have the options of either keeping my promise to you or breaking it. I know that five (equally weighty) promises made to you by others will only be kept if I break my promise. What should I do? An obvious answer is that I might ask you to cancel my promise (a power that you have as the promisee) so that you can benefit from your future promises being honored. However, suppose that this is not possible (perhaps you are unaware of the situation and I am unable to inform you). What should the agent do in such a circumstance? One possible answer is that the agent ought to break the promise because this will make the promisee better off overall. This answer involves a form of patient-relativity. However, some may have reservations about this answer. For, unlike the initial case where the agent was minimizing her own promise breaking, this case requires an agent to get her hands dirty by breaking a commitment that she otherwise would have kept, thereby committing a *prima facie* wrong. It might be argued that this is an unfair burden to place on the agent.¹⁵ Nonetheless,

15. This complaint may have even more intuitive force in the friendship example. For, unlike the good of promise keeping, the good of friendship may be equally good for both the patient and the agent and thus a duty to sacrifice friendship for the greater good of the friend may be an unfair burden on the agent.

some moral theories may endorse this kind of patient-relativity. Therefore, although patient-relativity may look most plausible in the case where it acts as a switch on moment-relativity, it could also act as a switch on agent-relativity, cancelling the agent-relative constraint when the trade-off involves the same patient.

This raises the question of how patient-relativity is related to the other distinctions we have discussed. In Section II, we saw that agent-relativity and moment-relativity are logically independent of one another and thus there are no limits on how agent-neutrality/relativity and moment-neutrality/relativity might be combined in a moral theory. Is the same true of patient-relativity? The talk above of its being a switch that turns on and off other kinds of relativity suggests that patient-relativity lacks logical independence. However, the switch metaphor is a little misleading here. For each of the cases we have considered so far involve not one patient-relative rule but rather two rules. For example, consider the initial promise-breaking example. One interpretation of the agent-relative constraint on breaking promises is moment-neutral. According to this interpretation the constraint on breaking promises requires agents to never break a promise unless doing so will lead to her breaking fewer promises overall. Such a rule can be succinctly stated as follows:

- (3) Each agent must minimize her own promise breaking.

Another interpretation of the agent-relative constraint on promise breaking is moment-relative. According to this interpretation the constraint on breaking promises requires agents to never break a promise in the present even when doing so will lead to less future promise breaking. This rule can be stated as follows:

- (4) Each agent, at each moment, must not break a promise.¹⁶

Finally, the patient-relative option requires agents to follow (3) in same-patient cases and follow (4) in different-patient cases. I submit that such a theory is best represented as containing two distinct rules:

- (5) Each agent must minimize her own promise breaking to the same patient.

16. You might worry that this statement does not adequately capture the moment-relative rule as it does not give guidance on what to do in situations where one will either break one promise now, or break several promises later, whereas a moment-relative rule properly conceived should instruct agents to favor present promise keeping. If you have this worry then the following may better capture the moment-relative rule: (4') Each agent, at each moment, must minimize her promise breaking at that moment.

(6) Each agent, at each moment, must not break a promise to one patient to bring about less promise breaking to other patients.

Rules (5) and (6) are each patient-relative because each rule makes reference to the numerical identity of moral patients. Rules (3) and (4) are patient-neutral because each rule makes no reference to the numerical identity of moral patients. Thus, we can say that a moral rule is patient-relative if and only if, when the rule is fully stated, it contains an ineliminable reference to the numerical identity of a moral patient, and is otherwise patient-neutral.

This intuitive account of the difference between a patient-neutral and a patient-relative rule can be adapted to the more formal approach employed by McNaughton and Rawling that was mentioned above in Section I. According to their approach a rule expressed with the form $(x)(x$ should ensure [. . .]) is agent-relative just in case there is an ineliminable occurrence of x in the square brackets that is bound by the initial universal quantifier. This approach can be adapted to moment-relativity by introducing a time variable “ t .” Thus, a rule expressed with the form $(x)(t)(x$ should ensure at t [. . .]) is moment-relative just in case there is an ineliminable occurrence of t in the square brackets that is bound by the initial temporal universal quantifier. Finally, we can say that a rule expressed with the form $(x)(x$ should ensure [. . .]) is patient-relative just in case there is an ineliminable occurrence of an identity operator in the square brackets (i.e., the operators “=” and “ \neq ”). For example, we might translate (6) above into McNaughton and Rawling’s canonical form as follows:

(6′) $(x)(t)(x$ should ensure at t [(y)(z) (If $y \neq z$ then it is not the case that x breaks a promise to y at t in order to keep more promises overall to z)]).

Notice that (6′) contains an ineliminable occurrence of an identity operator in the square brackets. Thus, according to the account of patient-relativity adapted to the McNaughton and Rawling approach, (6′) is correctly categorized as patient-relative.

Although it seems natural for a moral theory to combine rule (5) with rule (6), these two rules are logically independent as there is nothing contradictory in a moral theory containing one of these rules without containing the other. More generally, a moral theory can contain a rule proscribing an action in same-patient cases and yet contain no rule pertaining to what is permitted or prohibited in the corresponding different-patient cases (and vice versa). However, it seems reasonable to conjecture that any plausible moral theory containing a rule covering same-patient cases will also contain a rule covering different-patient cases (and vice versa). For, it would be strange for a moral theory to require one course of action

in same-patient cases if it didn't require a different course of action in different-patient cases. Thus, we can expect that, in the moral theories that interest us, patient-relative moral rules will generally come in pairs (covering same and different patient cases) whereas patient-neutral rules will not.

With this understanding of patient-relativity in place we can now see that there are various possible combinations of agent, moment, and patient-relativity/neutrality. The cases we have considered so far have only covered a few of these possible combinations. Table 1 summarizes all eight possible combinations.¹⁷

TABLE 1
POSSIBLE COMBINATIONS OF AGENT, MOMENT, AND
PATIENT-RELATIVITY/NEUTRALITY

Agent	Time	Patient	Example
1	N	N	Minimize general promise breaking
2	N	R	Minimize promise breaking to the same patient
3	N	N	Minimize present-moment promise breaking
4	N	R	Minimize present-moment promise breaking to the same patient
5	R	N	Minimize your promise breaking
6	R	R	Minimize your promise breaking to the same patient
7	R	N	Never break a promise, no matter the future consequences
8	R	R	Never break a promise to a patient, no matter the future consequences for that patient

NOTE.—N = neutral; R = relative.

It should be noted that any rule that combines agent and moment-relativity with patient-neutrality (i.e., combination 7) will entail a rule that combines relativity in all three domains (combination 8). For example, the rule "Never break a promise, no matter the future consequences" entails the rule "Never break a promise to a patient, no matter the future consequences for that patient" (meaning that the former rule prohibits everything that the latter rule prohibits).

No other entailments hold between the different combinations in this table. For example, it is not the case that rules that are neutral in every domain (combination 1) entail rules that are agent- and moment-

17. I thank Philip Pettit for urging me to consider all eight possible combinations summarized in this table.

neutral yet patient-relative (combination 2). One might think that such an entailment holds, but this is a mistake. For example, one might think that if you are required to minimize general promise breaking then you are required to minimize promise breaking to the same patient. However, this is not the case. In fact, these two rules are contraries. For the latter rule (combination 2) requires an agent to break a promise to a patient when that is the only way to bring about more promise keeping to that patient overall. However, this requirement conflicts with the former rule (combination 1) because there will be cases where breaking a promise will minimize promise breaking to the promisee and yet cause more promise breaking overall. In such cases the former rule requires promise keeping whereas the latter rule requires promise breaking.

C. Applications of Patient-Relativity

Above I motivated patient-relativity by presenting several cases where commonsense moral intuitions appear to require moral rules that make reference to the numerical identity of moral patients. However, we might ask whether there are moral theories that lend support to patient-relativity independently of these commonsense intuitions. In this section, I will consider three different moral theories that appear to provide a theoretical basis for patient-relativity. I will then show how patient-relativity provides a reply to Derek Parfit's "appeal to full relativity" argument.

First, consider the humanity formula of Kant's categorical imperative:

FH: Always treat people as ends in themselves and never as mere means.

The humanity formula is a very influential moral principle, and there are many different accounts of how to interpret it. On some, but not all, interpretations of this principle it gives the patient-relative verdict in the promise-keeping examples considered above. For breaking a promise now to one promisee to keep several promises in the future to other promisees might be classified as treating the present promisee as a "mere means" to fulfilling promises owed to several others. Furthermore, breaking a promise now to one promisee to keep several promises in the future to that same promisee might be classified as treating that promisee as an "end in herself" because it is out of concern for the promisee, qua rational agent, that the present promise is broken. It might be thought that this possibility demonstrates that we can account for the promise-keeping examples without appealing to patient-relative rules. For the humanity formula does not appear, at face value, to be a patient-relative rule (it makes no reference to the numerical identity of patients). However, I believe it shows the opposite. For the concepts used in the humanity formula of treating a person as a "mere means," and as an "end in herself,"

are imprecise and ambiguous in various ways. Indeed, the principle has many interpretations because different thinkers have taken these concepts to mean different things. Thus, when the principle is used in a particular example to derive a moral verdict we should ask: what interpretation of these key concepts is required to produce the claimed verdict? Perhaps there are several possible interpretations of “mere means” and “end in themselves” that allow the humanity formula to explain the promise-keeping examples. However, here is one interpretation worth considering:

To treat someone as a “mere means” is to perform an action with the motive of using that person to serve, or bring about, a goal that is not that person’s own goal but rather the goal of some numerically distinct person. To treat someone as an “end in themselves” is to perform an action with the motive of using that person to serve, or bring about, a goal that is that person’s own goal.

On this interpretation, the humanity formula turns out to be a patient-relative rule as, when fully stated, it makes reference to the numerical identity of moral patients. Several other obvious ways of interpreting the imprecise concepts in the humanity formula also produce a rule that is patient-relative. As such, I suspect that all plausible interpretations of the formula of humanity that allow it to explain the promise-keeping examples will be interpretations that make the principle patient-relative. However, arguing for this claim would be a long and difficult task and I will not attempt it here. Showing that, on some candidate interpretations of the formula of humanity, it is a kind of patient-relative rule that explains the promise-keeping examples is itself a significant conclusion. Thus, one theoretical basis for patient-relativity comes via certain interpretations of Kant’s humanity principle.

As a second example, consider the moral contractualism of T. M. Scanlon. According to Scanlon: “An act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behaviour that no one could reasonably reject.”¹⁸ Now, consider again the same-patient promise-keeping example. In this case there does not appear to be anyone who could reasonably reject the trade-off that is made on the patient-relative rule. For, there appears to be no one who can claim that they are being treated unfairly or unjustly. The patient involved has no reason to reject the trade-off and, in fact, has reasons to welcome it as she is better off in the long run by it (e.g., fewer promises owed to her are broken). The agent involved also has no reason to reject it and also appears to be better off as she

18. Thomas Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), 153.

minimizes her own wrongdoing. Finally, there appear to be no reasons for any third party to object to the trade-off being made. This can be contrasted with the “different-patient” case where some thing owed to one patient is withheld in order to give several other patients what they are owed. In this case the patients who benefit in the future, and the agent who performs the action, may be better off. However, the patient who loses in the present moment is worse off and can claim that she has been treated unfairly. On these grounds it could be reasonable to reject a principle that allowed this kind of trade-off. Thus, Scanlon’s contractualism can provide a theoretical grounding for patient-relativity.

As a third example consider the distinction between interpersonal aggregation and intrapersonal aggregation. Interpersonal aggregation is aggregation of the good across the lives of different patients. Intrapersonal aggregation is aggregation of the good in the life of a single patient. Some take this distinction to be morally significant.¹⁹ They typically endorse the following principle:

PA: It is morally permissible to harm a patient in order to maximize that patient’s total welfare. However, it is morally impermissible to harm a patient in order to maximize the total welfare of other patients.²⁰

PA is often motivated by appeal to the following considerations. It is normally rational for us to bear risks or accept certain harms to ourselves when doing so maximizes our overall welfare. Given this fact, it seems permissible for others to harm us in these ways in order to maximize our welfare. Yet many have the intuition that it is morally impermissible to harm someone to bring about a greater benefit to others. PA seems to adequately capture this set of intuitions.

PA is a patient-relative rule because it makes reference to the numerical identity of moral patients. Furthermore, PA entails the patient-relative verdict in the promise-keeping example. For making the trade-off in the different-patient case appears to be a form of interpersonal aggregation and thus is prohibited by PA. On the other hand, making the trade-off in the same-patient case appears to be a form of intrapersonal aggregation and thus is permitted by PA. However, although PA can

19. See Nagel, *The Possibility of Altruism*, 138; Scanlon, *What We Owe Each Other*, 229–40; and J. Paul Kelleher, “Relevance and Non-consequentialist Aggregation,” *Utilitas* 26 (2014): 386–408, 404–5. Nagel appeals to the separateness of persons to explain the moral significance of the distinction between intrapersonal and interpersonal aggregation.

20. Those who endorse this principle may qualify it in various ways. For example, they may agree that, above a certain threshold, harming in order to interpersonally aggregate is morally permitted. However, such a threshold is only applied to interpersonal aggregation; intrapersonal aggregation is taken to be permissible simpliciter.

explain patient-relativity in the promise-keeping case, it cannot explain all candidate cases of patient-relativity. For some moral rules that are unrelated to the aggregation of welfare can also be classed as patient-relative or patient-neutral. As an example, consider a rule that says the guilty ought to be punished. What does such a rule require in a case where the only options are: (i) punish one guilty person now, which will allow several people guilty of comparable crimes to escape punishment later; (ii) let one guilty person escape punishment now, which will ensure that several people guilty of comparable crimes get punished later? What about a case where the trade-off is between letting a guilty person get punished for one crime now but escape punishment for several comparable crimes later, or letting that person escape punishment for a crime now to ensure that she is punished for several later crimes? Some will endorse a rule that requires the trade-off in both kinds of cases because they believe we must maximize the punishment of the guilty. Such a rule is patient-neutral. Others will endorse a rule that prohibits the trade-off when there are different guilty patients but requires the trade-off with respect to the same guilty patient. Such a rule is patient-relative. Yet, on several prominent theories of retributive punishment, neither of these rules concerns the aggregation of patient welfare. Therefore, we might endorse patient-relativity in certain cases that cannot be explained by the principle PA.

It is also worth considering more generally how the distinction between interpersonal and intrapersonal aggregation relates to the patient-relative/patient-neutral distinction. Any moral rule concerned with interpersonal or intrapersonal aggregation will be a patient-relative rule as it will make reference to the numerical identity of moral patients. By contrast, any moral rule concerned only with aggregation simpliciter will be a patient-neutral rule as it will not make reference to the numerical identity of moral patients. An example of the former kind of rule is PA above. An example of the latter kind of rule is the aggregation rule endorsed by classical utilitarianism:

UA: The morally right action is that which maximizes total net welfare.

UA is an aggregation rule, yet in contrast to PA it is a patient-neutral rule as it makes no reference to the numerical identity of moral patients. It might be tempting to conclude from these observations that the patient-neutral/patient-relative distinction is equivalent to the distinction between moral rules that concern aggregation simpliciter and moral rules that concern interpersonal or intrapersonal aggregation. However, this is not the case, as moral rules that do not require any form of aggregation can also be classed as patient-relative or patient-neutral. For example, some deontologists advocate the “Equal greatest chance principle”

in rescue cases involving interpersonal trade-offs.²¹ They argue that if we find ourselves having to choose between saving the life of one stranger and saving the life of five strangers, and other than numbers all else is equal, then the morally correct thing to do is not to favor the larger number. Instead, we must give all strangers the greatest possible chance of being saved consistent with every stranger having an equal chance of being saved. In the case described we might do this by flipping a coin to decide whether to save the one or save the five, thereby giving each stranger a one in two chance of being saved. Deontologists who endorse this principle in rescue cases might also apply it to trade-offs involving promise breaking as follows:

EGC: When there are comparable promises owed by one promisor to different promisees, the promisor must give each promisee the greatest possible chance of having her promises honored, consistent with every promisee having an equal chance of having her promises honored.²²

EGC is a patient-relative rule because it makes reference to the numerical identity of moral patients. However, it is not a rule concerned with either interpersonal or intrapersonal aggregation. Thus, although all rules requiring interpersonal or intrapersonal aggregation are patient-relative, some patient-relative rules do not require either kind of aggregation. Therefore, the patient-relative/patient-neutral distinction is broader and more fundamental than the distinction between rules that require aggregation simpliciter and rules that require interpersonal or intrapersonal aggregation.

Finally, I will examine how recognizing patient-relativity can provide a response to Derek Parfit's "appeal to full relativity" argument. In chapter 9 of *Reasons and Persons*, Parfit develops an argument against rational egoism he calls "the appeal to full relativity." As we saw above in Section II, there are four possible ways a rule can combine agent and moment-neutrality/relativity. Two of these four possibilities combine relativity in one domain with neutrality in another domain. Parfit describes such rules as "incompletely relative." The other two possibilities were either neutral in both domains or relative in both domains. Parfit describes these respectively as "fully neutral" and "fully relative." According to Parfit, we should reject on theoretical grounds rules that are incompletely relative, and thus only accept rules that are fully neutral or fully relative. Parfit

21. For example, see John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6 (1977): 293–316.

22. EGC would typically be combined with a rule requiring promisors to minimize their promise breaking with respect to comparable promises owed to the same promisee.

claims this because he thinks there is an important analogy between “I” and “now”:

This claim can appeal to the analogy between oneself and the present, or what is referred to by the words ‘I’ and ‘now’. This analogy holds only at a formal level. Particular times do not resemble particular people. But the word ‘I’ refers to a particular person *in the same way in which* the word ‘now’ refers to a particular time. And when each of us is deciding what to do, he is asking, ‘What should *I* do *now*?’ Given the analogy between ‘I’ and ‘now’, a theory ought to give to both the same treatment.²³

Parfit directs the appeal to full relativity against the theory of rational egoism. However, he notes that this appeal is intended to apply to all normative theories and thus can be used to defeat any incompletely relative theory.²⁴ This has led some to argue that there are only two plausible kinds of moral theories, those that are agent- and moment-neutral and those that are agent- and moment-relative.²⁵ However, the most promising cases of patient-relativity appear to challenge this. For example, take the promise-breaking case discussed in part A. In this case we saw that there are strong intuitive and theoretical grounds for thinking that an agent-relative constraint on promise breaking must be moment-neutral with respect to the same patient. However, any moral rule that secured this result would be incompletely relative because it would combine agent-relativity with moment-neutrality. Thus, the intuitive and theoretical grounds in favor of this kind of patient-relativity are in conflict with Parfit’s theoretical considerations against incompletely relative rules. Given that the considerations in favor of patient-relativity appear strong and Parfit’s analogy argument appears weak, we should accept the possibility of incompletely relative moral theories. Importantly, once we accept an incompletely relative theory in any part of our normative theorizing, we no longer have grounds for rejecting other theories for their incomplete relativity. Thus, recognizing that there are strong arguments for patient-relativity not only demonstrates that there can be incompletely relative moral theories, it also undermines Parfit’s use of the appeal to full relativity against rational egoism.

IV. CONCLUSION

In this article, I have presented a new kind of moral relativity that stands alongside agent-relativity and moment-relativity. All moral rules can be

23. Parfit, *Reasons and Persons*, 140.

24. *Ibid.*, 148.

25. Louise, “Relativity of Value,” 535.

classified as either agent-relative or agent-neutral. They can also be classified as either moment-relative or moment-neutral. This article shows that there is a third type of classification that needs to be added to this list. A moral rule can also be classified as either patient-relative or patient-neutral. Being aware of this third kind of relativity improves our understanding of the main structural features of moral theories. It also allows us to better understand how agent-relativity and moment-relativity interact. Finally, it appears to give us plausible examples of normative theories that are incompletely relative and thus provides us with a response to Derek Parfit's "appeal to full relativity" argument.